

THE MULTI-OBJECT, FIBER-FED SPECTROGRAPHS FOR SDSS AND THE BARYON OSCILLATION SPECTROSCOPIC SURVEY

STEPHEN A. SMEE¹, JAMES E. GUNN², ALAN UOMOTO³, NATALIE ROE⁴, DAVID SCHLEGEL⁴, CONSTANCE M. ROCKOSI⁵, MICHAEL A. CARR², FRENCH LEGER⁶, KYLE S. DAWSON⁷, MATTHEW D. OLMSTEAD⁷, JON BRINKMANN⁸, RUSSELL OWEN⁶, ROBERT H. BARKHOUSER¹, KLAUS HONSCHIED⁹, PAUL HARDING¹⁰, DAN LONG⁸, ROBERT H. LUPTON², CRAIG LOOMIS², LAUREN ANDERSON⁶, JAMES ANNIS¹¹, MARIANGELA BERNARDI¹², VAISHALI BHARDWAJ⁶, DMITRY BIZYAIEV⁸, ADAM S. BOLTON⁷, HOWARD BREWINGTON⁸, JOHN W. BRIGGS¹³, SCOTT BURLES¹⁴, JAMES G. BURNS⁹, FRANCISCO JAVIER CASTANDER^{15,16}, ANDREW CONNOLLY⁶, JAMES R. A. DAVENPORT⁶, GARRETT EBELKE⁸, HARLAND EPPS⁵, PAUL D. FELDMAN¹, SCOTT FRIEDMAN¹⁶, JOSHUA FRIEMAN¹¹, TIMOTHY HECKMAN¹, CHARLES L. HULL³, GILLIAN R. KNAPP², DAVID M. LAWRENCE⁷, JON LOVEDAY¹⁷, EDWARD J. MANNERY⁶, ELENA MALANUSHENKO⁸, VIKTOR MALANUSHENKO⁸, ARONNE JAMES MERRELLI¹⁸, DEMITRI MUNA¹⁹, PETER R. NEWMAN⁸, ROBERT C. NICHOL²⁰, DANIEL ORAVETZ⁸, KAIKE PAN⁸, ADRIAN C. POPE²¹, PAUL G. RICKETTS⁷, ALAINA SHELDEN⁸, DALE SANDFORD⁵, WALTER SIEGMUND⁶, AUDREY SIMMONS⁸, D. SHANE SMITH⁹, STEPHANIE SNEDDEN⁸, DONALD P. SCHNEIDER^{22,23}, MICHAEL STRAUSS², MARK SUBBARAO¹², CHRISTY TREMONTI²⁴, PATRICK WADDELL²⁵, DONALD G. YORK²⁶

Draft version August 13, 2012

ABSTRACT

We present the design and performance of the multi-object fiber spectrographs for the Sloan Digital Sky Survey (SDSS) and their upgrade for the Baryon Oscillation Spectroscopic Survey (BOSS). Originally commissioned in Fall 1999 on the 2.5-m aperture Sloan Telescope at Apache Point Observatory, the spectrographs produced more than 1.5 million spectra for the SDSS and SDSS-II surveys, enabling a wide variety of Galactic and extra-galactic science including the first observation of baryon acoustic oscillations in 2005. The spectrographs were upgraded in 2009 and are currently in use for BOSS, the flagship survey of the third-generation SDSS-III project. BOSS will measure redshifts of 1.35 million massive galaxies to redshift 0.7 and Lyman- α absorption of 160,000 high redshift quasars over 10,000 square degrees of sky, making percent level measurements of the absolute cosmic distance scale of the Universe and placing tight constraints on the equation of state of dark energy.

The twin multi-object fiber spectrographs utilize a simple optical layout with reflective collimators, all-refractive cameras, gratings, and state-of-the-art CCD detectors to produce hundreds of spectra simultaneously in two channels over a bandpass covering the near ultraviolet to the near infrared, with a resolving power $R = \lambda/\text{FWHM} \sim 2000$. Building on proven heritage, the spectrographs were upgraded for BOSS with volume-phase holographic gratings and modern CCD detectors, improving the peak throughput by nearly a factor of two, extending the bandpass to cover $360 < \lambda < 1000$ nm, and increasing the number of fibers from 640 to 1000 per exposure. In this paper we describe the original SDSS spectrograph design and the upgrades implemented for BOSS, and document the predicted and measured performances.

Subject headings: instrumentation: spectrographs

smee@pha.jhu.edu

¹ Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

² Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

³ Observatories of the Carnegie Institution of Washington, 813 Santa Barbara Street, Pasadena, CA 91101, USA

⁴ Physics Division, Lawrence Berkeley National Laboratory, Berkeley CA 94720, USA

⁵ UC Observatories and Department of Astronomy and Astrophysics, University of California, Santa Cruz, 375 Interdisciplinary Sciences Building (ISB) Santa Cruz, CA 95064, USA

⁶ Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA

⁷ Department of Physics and Astronomy, University of Utah, Salt Lake City, UT 84112, USA

⁸ Apache Point Observatory, Sunspot, NM 88349, USA

⁹ Department of Physics and Center for Cosmology and Astroparticle Physics, Ohio State University, Columbus, OH 43210, USA

¹⁰ Department of Astronomy, Case Western Reserve University, Cleveland, OH 44106, USA

¹¹ Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA

¹² Department of Physics and Astronomy, The University of Pennsylvania, 209 South 33rd Street, Philadelphia, PA 19104, USA

USA

¹³ HUT Observatory, Mittelman Family Foundation, P.O. Box 5320, Eagle, CO 81631, USA

¹⁴ Physics Department, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

¹⁵ Institut de Ciències de l'Espai (IEEC-CSIC), E-08193 Bellaterra, Barcelona, Spain

¹⁶ Space Telescope Science Institute, Baltimore, MD 21218, USA

¹⁷ Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QJ, UK

¹⁸ Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

¹⁹ Center for Cosmology and Particle Physics, New York University, 4 Washington Place, New York, NY 10003, USA

²⁰ Institute of Cosmology and Gravitation (ICG), Dennis Sciana Building, University of Portsmouth, Portsmouth, PO1 3FX, UK

²¹ High Energy Physics Division, Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA

²² Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

²³ Institute for Gravitation and the Cosmos, The Pennsylvania State University, PA 16802, USA

²⁴ Department of Astronomy, University of Wisconsin-

1. INTRODUCTION

The Sloan Digital Sky Survey (SDSS; York et al. 2000) project was conceived in the mid-1980s as an ambitious endeavor to understand the large-scale structure of the Universe. SDSS and its extension, SDSS-II, conducted a coordinated imaging and spectroscopic survey from 2000-2008 over approximately $10,000 \text{ deg}^2$ of high Galactic latitude sky. Now in its third phase of operation, SDSS is one of the most successful projects in the history of astronomy. The survey has produced an enormous catalog consisting of five-band digital images that include nearly one billion unique objects, and spectra of 930,000 galaxies, 120,000 quasars, and 460,000 stars, all publicly available, (Abazajian et al. 2009, and references therein).

To obtain these imaging and spectroscopic data, a dedicated 2.5m telescope (Gunn et al. 2006), wide-field mosaic CCD camera (Gunn et al. 1998), and twin multi-object fiber spectrographs were constructed and installed at the Apache Point Observatory (APO), in Sunspot, New Mexico. The telescope, built to accommodate the requirements for both imaging and spectroscopy, is shared by the camera and spectrographs, which mount at the Cassegrain focus. The imaging survey was carried out on clear, dark nights with good seeing using the 120 mega-pixel camera, which operated in drift-scanning mode using a 5×6 array of 2048×2048 pixel detectors to obtain *ugriz* (Fukugita et al. 1996), photometry. The imaging data, once reduced and calibrated (Smith et al. 2002; Pier et al. 2003; Ivezić et al. 2004; Tucker et al. 2006; Padmanabhan et al. 2008), were used for spectroscopic target selection. Spectroscopy was performed using the two multi-object fiber spectrographs, collecting 640 spectra over the 3° diameter field in one exposure.

In this paper, we describe the design and performance of the SDSS spectrographs, and their recent upgrade for the Baryon Oscillation Spectroscopic Survey (BOSS Schlegel et al. 2009; Dawson et al. 2012). BOSS is the flagship survey in the third-generation SDSS-III program currently underway at the 2.5m SDSS telescope (Eisenstein et al. 2011). BOSS will measure the cosmic expansion history of the universe to percent-level precision by mapping an immense volume of sky to obtain the spatial distributions of galaxies and quasars, and from it, the characteristic scale imprinted by baryon acoustic oscillations (BAO) in the early universe (for a review of BAO with a respect to other cosmological probes, see Weinberg et al. 2012). A measure of the scale at low redshifts, out to $z \sim 0.7$, will be obtained by carrying out a redshift survey of 1.35 million massive galaxies from $10,000 \text{ deg}^2$ of SDSS data. BOSS will also observe Lyman- α absorption in the spectra of 160,000 high-redshift quasars to measure large scale structure at redshifts of $z \sim 2.5$.

The SDSS spectrographs utilize a dual-channel design with a common reflecting collimator and a dichroic to split the beam into a blue channel and a red channel. In each channel, just downstream of the dichroic, a transmitting grating disperses the light, which is imaged by an all-refractive camera onto a CCD. For BOSS the basic

optical design has been retained, with several improvements. The ruled gratings have been replaced by volume-phase holographic (VPH) gratings (gratings sandwiched between two prisms) and the CCDs have been replaced with more modern devices. These changes produce a significant improvement in throughput and a modest extension of the wavelength range in both the blue and red channels. Additionally, smaller diameter fibers that are better matched to the angular scale of BOSS targets have been installed, allowing the total number of simultaneous spectra obtained from the two spectrographs to be increased from 640 in the original design to 1000 in the BOSS configuration.

The remainder of this paper is organized as follows. In Section 2 we begin by describing the design and construction of the original SDSS spectrographs in some detail, published here for the first time. This is followed in Section 3 by a discussion of the spectrograph upgrades completed in 2009 for BOSS. The performance of both the original SDSS spectrographs and the upgraded BOSS design is presented in Section 4. Finally, a summary of the spectrograph design and performance is provided in Section 5 along with some highlights of the scientific research enabled by these instruments.

2. SDSS SPECTROGRAPH DESIGN

2.1. Design Requirements

The requirements for the SDSS spectrographs were set by its primary scientific goal: the creation of a three-dimensional wide-area map of the universe to reveal its large-scale structure. The SDSS imaging survey provides the two-dimensional locations of nearly one billion celestial objects, and spectroscopy of a selected subset of targets is then used to determine redshifts and thus distances. The project set as a requirement spectroscopy of one million galaxies and 100,000 quasars distributed over approximately $10,000 \text{ deg}^2$.

Acquisition of a large number of spectra simultaneously over a large field of view, with moderate resolution sufficient for accurate redshift measurements, naturally led to the choice of a fiber-fed multi-object spectrograph. The spectrograph design was also dictated in large part by the design of the telescope, which was itself optimized for both wide-field, multi-band, imaging and multi-object spectroscopy. In what follows we summarize the requirements which dictated the design of the SDSS spectrographs.

2.1.1. Telescope Design

The 2.5m SDSS telescope is a modified distortion-free Richey-Chrétien design with a 3° diameter field of view and $f/5$ final focal ratio. The optical design incorporates two aspheric corrector lenses, a Gascoigne-type design located near the vertex of the primary mirror, and two interchangeable secondary correctors, one used for imaging and the other for spectroscopy. The imaging corrector is a thick fused silica lens located close to the focal plane that is incorporated into the SDSS camera, where it serves a mechanical function in addition to providing optical correction. The spectroscopic corrector is a thinner lens located further from the focal plane and optimized for chromatic focus. The plate scale for spectroscopy is $3.627 \text{ mm/arcminute}$. The spectroscopic focal surface is

Madison, Madison, WI 53703, USA

²⁵ NASA Ames Research Center, Moffett Field, CA 94035, USA

²⁶ Department of Astronomy and Astrophysics and the Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

slightly curved, with a maximum deviation from a plane of 2.6 mm. One important detail of the spectroscopic optics is that the central ray is not perpendicular to the focal plane, necessitating a clever correction scheme for fiber placement that will be described below.

2.1.2. Number of Fibers

Spectroscopy of approximately one million objects over $10,000 \text{ deg}^2$, plus 10 – 20% additional fibers for calibration sources and sky background measurements, implies a density of 120 deg^{-2} . The 2.5m telescope has a field of view of 7 deg^2 , but each plate will view a unique area on the sky of about 5 deg^2 . The higher density of plates is due to the need for overlap between fields to ensure complete sky coverage without gaps, and to allow multiple observations to cross-calibrate the entire survey. The required number of fibers is therefore approximately 600 per plate.

A practical limit on the number of fibers was imposed by the detector format, camera design, and fiber mounting scheme. For proper spectral sampling, the fiber images on the detector should be about 3 pixels in diameter, with an equal space between spectra to reduce crosstalk and allow for a measurement of the scattered light floor. Thus, each spectrum used six detector columns, and the 2048×2028 pixel detector could accommodate a maximum of 341 spectra. The actual number was reduced to 320 spectra to avoid camera optical distortions near the detector edges and to allow for extra gaps between groups of 20 fibers, which was necessary for the fiber mounting scheme described in Section 2.2. These larger gaps turned out to be quite useful for measurements of scattered light in the wings. The final choice of 640 fibers per plate, or 320 per spectrograph, provided some contingency over the required 600 fibers, allowing for broken fibers, additional calibration fibers and/or ancillary programs utilizing the extra fibers.

2.1.3. Fiber Diameter

The fiber diameter is set by the desire to maximize the signal-to-noise (S/N) ratio for an extended source given the sky background. For the galaxies of interest around redshift $z = 0.1$ and the sky conditions at Apache Point, this corresponds to a fiber size of around $3''$, or a fiber diameter of 180 microns. Fibers of good optical quality were also readily obtainable in this size.

2.1.4. Resolving Power

The spectroscopic resolution is defined as the full-width at half maximum (FWHM) of the one-dimensional point spread function (PSF), in wavelength units (a *resolution element*). The resolving power is the wavelength divided by this quantity, and we will often use the phrase “higher resolution” to mean higher *resolving power*, as is the normal usage. Given a fixed number of pixels in the dispersion direction and requiring proper sampling, increasing resolving power reduces the wavelength range. Higher resolving power also reduces the number of source photons per pixel, increasing the exposure time required to exceed the CCD read noise. On the other hand, if the resolving power is too low, absorption lines cannot be resolved and this will ultimately degrade both the accuracy and success rate of redshift measurements.

The resolution was therefore set by the requirement that broadened absorption lines be well resolved. Although there is a large range in galaxy velocity dispersions, a typical galaxy has a velocity dispersion of about 100 to 200 km/s, corresponding to a resolving power of 1500-3000. The actual resolving power as a function of wavelength was allowed to vary within these limits to optimize the red-blue channel wavelength split location and the total wavelength coverage, while maintaining well-sampled spectra with 3 pixels per resolution element on the CCD over the full wavelength range.

2.1.5. Wavelength Range

Redshifts are determined either from absorption lines or emission lines – in both cases only a few lines contribute most of the signal. In absorption, three features are dominant: the Mg b triplet at $\lambda = 5180 \text{ \AA}$, Ca at $\lambda = 5270 \text{ \AA}$, and the Na I doublet (D lines) at $\lambda = 5890 \text{ \AA}$. At shorter wavelengths, the Ca II K and H lines at $\lambda = 3933, 3969 \text{ \AA}$ and the G band 4300 \AA may also be detected in absorption. In emission, $H\alpha = 6563 \text{ \AA}$ is the strongest (and often the only) line, although the [OII] doublet may also be visible at $\lambda = 3727 \text{ \AA}$.

Given the availability of these spectral features, and considering practical limitations on UV throughput, the short wavelength cutoff was set at 3900 \AA to ensure that the H and K lines of CaII are observable even at zero redshift, while the [OII] doublet is observable at $z > 0.05$. Redshift determination for most nearby galaxies could have been accomplished with a single blue arm extending up to 6000 \AA ; however, the SDSS imaging camera was designed to measure to the detector red limit cutoff, so it was decided to take advantage of the detector sensitivity in the spectrographs and add the red channel. This would enable observation of H- α to a redshift of $z = 0.2$ or more, as well as the observation of quasars out to redshifts beyond $z = 5$.

The extension of upper wavelength cutoff above 8000 \AA opened up a rich new vein of scientific discovery that was not anticipated at the time of the instrument design. In particular, pushing the long wavelength cutoff as high as possible extended the limit for redshift determination of luminous red galaxies (LRG’s) using the 4000 \AA break. The LRG sample (Eisenstein et al. 2001) was used to make the first observation of the baryon acoustic oscillation feature, which in turn motivated the future upgrade of the spectrographs to even longer wavelengths for BOSS, as discussed later in this paper.

2.1.6. Throughput/Signal to Noise

The requirement on throughput was set by the desire to obtain over one million spectra over $10,000 \text{ deg}^2$ to a limiting Petrosian magnitude of $r = 18.15$ in five years, corresponding to roughly 100 deg^{-2} galaxies. Given the average seeing and extinction at APO, this implied a maximum exposure time of one hour. The corresponding throughput requirement, including atmospheric extinction and the telescope throughput, varies as a function of wavelength and has a maximum of about 17% at 7000 \AA .

Provided that the spectral resolution is sufficient to resolve the absorption lines, the minimum S/N ratio

needed to derive a redshift depends mainly on the strength of the absorption lines. For convenience, the S/N per \AA of spectral continuum will be quoted. For an elliptical galaxy with strong lines, spectra obtained in the Center for Astrophysics redshift surveys (Huchra et al. 1983; Falco et al. 1999) show that one can measure a reliable redshift if S/N per \AA > 8 , i.e., one needs to collect 64 object photons per \AA , assuming that the noise is dominated by photon statistics from the source. This number must be increased, however, if sky background and/or readout noise is significant. A significant problem for some galaxies is that they have weak absorption lines (presumably because they have a significant amount of light from early-type stars) and yet lack strong $H\alpha$ emission. In these cases one may need two or three times as many photons to derive an absorption-line redshift. We adopt as a guide the goal of obtaining spectra with S/N of 15 per \AA . Simulated galaxy and quasar spectra indicated that we could in fact reach this goal with exposures of somewhat less than a one hour.

2.2. Fiber System Design

2.2.1. Overview of Fiber System

Light is transmitted from the telescope focal plane to two identical spectrographs by fiber optic strands 180 microns in diameter ($3''$ on the sky). Light enters the fibers at the telescope focal plane in a cone of numerical aperture 0.1 (f/5 beam), and the spectrograph collects light emitted from the other end of the fibers in a slightly larger cone with numerical aperture 0.25 (f/4), due to focal ratio degradation (FRD) that occurs as the light travels down the fiber. Any light emitted outside this cone will be lost, so a primary requirement on the fiber system is to limit FRD so as to maximize throughput. To this end, the spectrographs are mounted on the telescope to avoid any relative motion between the two ends of the fibers and the resulting stress that can result in increased FRD. This scheme also maximizes throughput by keeping the fibers short, minimizes fiber throughput variations due to physical motion and stress, and avoids the problems of routing and protecting long fiber runs. The sky ends of the fibers are plugged into drilled 0.813 m diameter aluminum plates called *plug plates* that position the fibers on the spectrograph focal plane, and the other ends of the fibers are terminated in one of two slitplates. Each thin slitplate is mounted to a rigid frame with precision locating features for accurate placement in the spectrograph. The assembly of plug plate, fibers and slitheads is mechanically supported by a portable aluminum cartridge that can be installed on the telescope by a single operator in a few minutes. New plug plates are mounted on the cartridges during the day and plugged with fibers, then sequentially mounted on the telescope during the night. A rendering and photograph of a fiber cartridge are shown in Figures 1 and 2, respectively.

For each new sky field, a cartridge is wheeled under the telescope using the Linde cart (named for its designer Carl Lindenmeyer) and attached to the telescope rotator using pneumatic clamps. At the same time the attached slitheads enter the spectrographs through the open slithead doors and are clamped in place. A kinematic mounting interface ensures accurate, repeatable placement. The photograph in Figure 3 illustrates the

operation. Eight cartridges were fabricated for SDSS to provide sufficient pre-plugged plates for an entire night of observing.

Each cartridge also has a set of coherent fibers that are placed on pre-selected guide stars and viewed by the guider camera. These stars are used for field rotation and translation at setup time.

2.2.2. Cartridges

The fiber cartridge consists of a machined aluminum cast body that supports the optical fiber harnesses, spectrograph slitheads, and plug-plate. Assembling these components into a single robust unit protects the fragile fibers during the manipulations necessary for plugging, transport to and from the telescope, and mounting onto the instrument rotator. The cartridges are plugged during the day, and are designed so they can be quickly installed on the telescope at night under often difficult conditions of low light and cold temperatures.

The plug-plate holder consists of two large bending rings that warp the plug-plate to match the telescope best-focus surface. The bending rings are mounted to the cartridge body with a set of kinematic pin mounts. The alignment of the cartridge to the telescope is provided by another kinematic mount employing v-groove blocks on the cartridge that engage with v-blocks on the telescope. This system ensures proper and repeatable alignment between the plug plate and telescope focal plane.

The two slitplates, each supporting 320 fibers, are mounted in their respective slitheads. These slitheads are aluminum assemblies mounted outboard of the cartridge body that support and protect the slitplates. The slitheads are attached to the cartridge body with a spring-loaded seating system that provides alignment for insertion into the spectrograph bodies, but then allows the slithead to float free from the cartridge body and engage the slithead-to-spectrograph kinematic mounting system. When not mounted on the telescope, these slitheads are protected by sliding covers to prevent contamination and/or mechanical contact with the delicate slitplates.

All cartridge operations occur at the same elevation, on the telescope platform and the adjacent support building. In the plugging lab, the exposed plug-plates from the previous night's observing are unplugged and removed from their cartridges and new plates are installed. Once plugging and fiber mapping is completed (a process that takes 30 – 45 minutes), the 145 kg cartridge is stowed on a lift table installed in a bay that provides both interior and exterior bay door access. At night, the outside door is opened to allow the cartridges to equilibrate to the temperature of the ambient air.

To install a new cartridge on the telescope, an outside manipulator arm is employed to move the cartridge from the storage bay to one of *two* receiving plates on the Linde cart. The cartridge is then wheeled from the storage bay to the telescope. With the telescope parked at zenith and locked into position, the Linde cart is rolled under the mounted cartridge to align the *empty* receiver plate with it. Aided by a hydraulic lift, the observer removes the exposed cartridge from the telescope. Then, maneuvering the Linde cart to align the unexposed cartridge onto the hydraulic lift, the observer mounts the new cartridge onto the telescope instrument




FIG. 1.— Rendering of a fiber cartridge. The fiber cartridge consists of a cast aluminum body that supports the fiber harness, the two slitheads, and the plug-plate. The slitheads are attached to the cartridge body with a spring-loaded seating system that provides alignment for insertion into the spectrograph bodies, but then allows the slithead to float free from the cartridge body and engage the slithead-to-spectrograph kinematic mounting system. Kinematic mounts around the periphery of the cartridge casting ensure accurate and repeatable placement of the cartridge with respect to the telescope.




FIG. 2.— Photograph of a BOSS fiber cartridge. Fibers plugged into the back of the plug-plate are routed in bundles to the slitheads (the two boxes standing upright at the left and right side of the cartridge). The design shown is identical to that used for SDSS except for the number and size of the fibers. For SDSS, 320 fibers are routed to each slithead, while for BOSS each slithead carries 500 fibers.

rotator. Once the new cartridge is latched and the cart receiver plate is lowered back onto the Linde cart, the cart is rolled out from under the telescope. The telescope is now ready to move to the next field and to begin an-




FIG. 3.— Photograph showing a fiber cartridge being installed on the telescope. The cartridge is raised by a hydraulic lift in the floor below the primary cell. When raised, the cartridge engages kinematic mounts for precise location. At the same time, the two slitheads engage the spectrographs, each of which is located by its own kinematic mounting features integral to the slithead and spectrograph optical bench. Installation takes approximately three to five minutes.

other exposure. Only three to five minutes are required to perform this cartridge change.

As the cartridge is lifted into place and clamped to the telescope, the slitheads are simultaneously inserted into sockets in the spectrographs. The slitheads are at-

tached to the cartridge frame by stiff springs so that they can move slightly with respect to the rest of the cartridge. Once the cartridge has been correctly positioned and clamped to the telescope, each slithead is loaded against a three-point kinematic mount on the spectrograph by a single pneumatic clamp. A flexible rubber seal between the slit heads and the spectrograph bodies prevents extraneous light from entering during exposures. Each slithead is coded and its identification relayed to the observer's workstation when it is inserted. This information allows adjustments for each slithead, e.g., image placement on the CCD and focus, to be made automatically. Figure 4 shows two schematic views of the cartridge mounted on the telescope with the slitheads inserted into the spectrographs.




FIG. 4.— Two schematic views of the cartridge mounted on the telescope. Top: a cutaway side view showing the slitheads inserted into the spectrographs (only 9 fiber harnesses are shown). Bottom: top view showing the cartridge located between the two spectrographs which are mounted to the instrument rotator (depicted as the large outer circle).

2.2.3. Optical Fiber

The selected optical fiber material is a silica UV-enhanced step-index fiber with a core diameter of 180 microns, a thin cladding and a polyimide protective layer. The actual fiber was Polymicro Technologies, Inc. FHP 180-198-218, where the numbers refer to the diameter of the bare fiber, plus cladding and plus polyimide buffer.

2.2.4. Fiber Harnesses

A fiber harness consists of 20 fibers of length 1.865 ± 0.025 meters; each fiber cartridge contains 32 fiber harnesses. The fiber harness is the unit procured from the vendor responsible for preparing the fibers by terminating both ends. A fiber harness is shown in the photograph in Figure 5.



FIG. 5.— Fiber harness for the SDSS spectrographs. The ends, which are plugged and unplugged during operations, are protected by tough nylon tubing, which terminates in the anchor block. The lens enlarges the v-groove block termination at the slit end.

The plugging end is terminated by gluing it into a ferrule and polishing the end. The ferrule is machined from stainless steel; the most important tolerances on the ferrule are the centering and angle of the hole in which the fiber is inserted, specified at 12 microns and 1.1° , respectively. The tilt of the plane of the polish on the fiber end is specified at less than 1.0° . The fiber is glued into the ferrule using a low-shrinkage adhesive selected to retain flexibility for many years over a temperature range -20 to $+30^\circ$ C.

To facilitate handling of the fibers during plugging and unplugging while protecting the fragile fibers, a protective jacket 38 cm long is attached to the ferrule on one end and glued into an anchor block on the other end. The selected jacket material consists of nylon tubing 3.2 mm in diameter, allowing the fiber to move freely inside to avoid imparting any stress. The anchor block holds twenty fiber jackets in a row on 5.1 mm centers. The anchor blocks are attached to the cartridge body, providing a fixed point from which the fibers are routed to the slitplate. The fibers are stationary from the anchor block to the slitplate, and the only portion of the fiber which moves during plugging is protected inside the nylon jacket. Thin tubing, 95 cm long and 2 mm in diameter, is used to bundle the fibers into groups of 10 immediately before their termination at the slitplate.

At the slitplate end, twenty fibers at a time are mass-terminated by gluing them into a machined piece of stain-

less steel called the v-groove block; see Figure 6. This item is so named because it has twenty v-shaped grooves machined on an electrical discharge machine (EDM) to precisely locate the fibers. The grooves are not parallel, but fan out slightly so they are normal to the tangent of a circle with radius 640 mm. For termination, twenty fibers are arranged side by side in the v-groove block on 390 μm centers and glued into place with a cover plate on top. The tolerance on the fiber placement is 30 μm (individual, not cumulative), and 0.3° in tilt between the groove and length of the fiber; there is an additional tolerance of 0.3° on the alignment of the grooves with respect to the body of the v-groove block. After the fibers are aligned in the v-groove block, covered and glued in place, the ends are polished on a flat surface, with a specified tolerance of 1.0° on the tilt of the plane of polish.




FIG. 6.— Three views of the v-groove blocks as mounted on the slit plate. Each block contains 20 fibers, which are aligned by 20 v-shaped grooves. The fibers are held in place by epoxy sandwiched between the fibers and a cover plate. A hole in the slit plate under each block permits easy removal for replacement.

2.2.5. Slitplates

Each spectrograph accepts the light from 320 fibers, terminated in 16 v-groove blocks that are glued to a thin, rigid aluminum slitplate with a radius of curvature of 640 mm. The center-to-center spacing between fibers in adjacent v-groove blocks is 523 microns, compared to 390 microns between fibers within a v-groove block. The v-groove block is aligned by pressing the output edge against a curved alignment jig, and attached to the slit plate with an adhesive. An air-powered fluid dispenser meters the adhesive to prevent contamination of critical surfaces. The slitplates are in the beam and therefore must be as thin as possible. The total thickness of the plate and the v-groove blocks is specified to be less than 3.2 mm.

2.2.6. Plug Plates

The telescope optical system is a simple, fast, large field design with a focal surface flat to 2.6 mm but one where the principal ray deviates from the normal to the

best-focus surface by up to 37 milliradians (2.12°). For highest efficiency, the ends of the optical fibers should be positioned on the best-focus surface with their axes aligned with the principal ray. Plug-plate technology can be made to satisfy these criteria quite nicely.

The plug-plates are an aluminum alloy 2024-T3, 3.2 mm thick and 0.813 m in diameter. By applying bending moments to the edge of the plate (beyond the field of view), finite element calculations show they can be deformed to match the best-focus surface to an area-weighted 62 μm RMS. The greatest departure from the best-focus surface is 200 microns and occurs at the center where the images are the best. Overall, the images are not significantly degraded from the best-focus surface.

As deformed to match the best-focus surface, the hole axes should align with the principal ray axes. This configuration is straightforward to accomplish if the plug-plate is deformed (in the opposite sense) over a properly curved mandrel, for drilling. The drilling is performed using a three-axis Computer Numerically Controlled (CNC) milling machine, as it is not necessary to tilt the drilling head or the plug-plate.

Drill test results indicate that holes can be drilled with an accuracy of 9 μm RMS in position and 4 μm RMS in diameter using short high-precision spade drill bits in a custom-made collet. In the test, four different bits were used to drill 50 holes each. The drilling time was 5.8 sec/hole. No significant degradation in drilling accuracy was observed for a range of slopes in the work-piece surface from 0 to 70 milliradians.

The plug-plates have a mass of 4.3 kg and are sufficiently thin so that the bending stresses, forces and material costs are reasonable. They are thick enough to provide hole depth adequate to constrain the plug angular alignment with the hole and to prevent significant gravity-induced deflections.

Once plugged, profilometry is performed on the plug plates in the plugging laboratory to verify that the surface conforms to the focal plane within tolerances. This is done using a fixture that mounts five digital linear indicators on a steel beam section that straddles the plug plate, using a kinematic support system to make contact with the aluminum. The five indicators measure the plate displacement from a calibrated 0.000 mm baseline along a radial direction from the center to the edge of the plug plate. There are four sets of measurements taken, spaced 90° apart for each plug plate. The measurements are then recorded in a database for each plugging of each plate.

2.2.7. Fiber Tester

Each completed fiber harness is tested by the vendor to ensure that each of the twenty fibers meets a minimum throughput requirement of 85%. The test device, which was supplied to the vendor, uses white light from an intensity-stabilized quartz-halogen lamp which is imaged onto the fiber under test using a source fiber and a microscope objective, which produces a uniform f/5 converging beam. A microscope eyepiece and pellicle beam-splitter provide a view of the fiber under test.

Light from the output end of the test fiber is collected by a pair of achromatic doublets focused onto a silicon photodiode. A filter between the doublets flattens the quartz-halogen spectral curve. A calibrated aperture

blocks light outside a cone of $f/4$. A computer-controlled translation stage accurately locates the appropriate fiber of the v-groove block in front of the aperture. This same light collection system is also used to measure light from the microscope objective, in order to make absolute throughput measurements.

An identical instrument was used to verify the manufacturer's measurements for some fraction of the fibers.

2.2.8. Fiber Mapper

With 640 fibers on each plug plate, it was important not to rely on manual plugging of each fiber into a specific pre-assigned hole. However, plugging cannot be completely random because not every fiber can reach every position on the plate; this constraint was adopted to keep the total fiber length relatively short. The solution was to have hand-marked regions on each plate for specific fiber sets that can reach all of the holes within the delineated area, and to then allow for random plugging within those regions. After plugging is complete, the cartridge is wheeled over to a fiber mapper system which is used to determine which fiber is in which hole.

The fiber mapper illuminates each fiber sequentially by moving a light along the slit head; the illuminated fiber appears as a bright point against the dark background of the plug plate. A CCTV camera behind a narrow band filter matching the light source determines the X-Y position of each fiber on the plugplate. This X-Y position is converted to sky coordinates and stored in the database, providing a map between the physical location of each spectra on the CCD and the corresponding coordinates on the sky. The fiber mapper is also useful for identifying broken or loose fibers that drop down into the hole after plugging. The entire operation takes about five minutes.

2.3. Optical Design

2.3.1. Optical Design Overview

The optical design of the spectrographs was strongly influenced by the use of optical fibers to feed light from the telescope to the spectrographs. With a scale of $16.7''/\text{mm}$ at the telescope focal plane, the required $3''$ entrance aperture corresponds to a fiber diameter of 180 microns. The $f/5$ focal ratio of the telescope matches well to the fiber acceptance angle and results in an $f/4$ exit beam, expanded somewhat by the FRD scattering processes within the fibers.

With the fiber interface so defined, the basic optical design of the spectrographs is rather straightforward. The collimator must operate at $f/4$ or faster, or some of the light exiting the fibers will be lost. The exit face of a fiber, which is the entrance aperture into the spectrograph for one object, must be re-imaged onto the detector with adequate pixel sampling. A camera focal ratio of $f/1.5$ results in a demagnification of 2.67 and leads to fiber images that are just under three pixels on the $24\text{ }\mu\text{m}$ pixel pitch, SITe $2\text{ k} \times 2\text{ k}$ format CCDs. At this sampling, approximately 700 resolution elements are available across the detector. Clearly, the full 3900-9100 Å bandpass will not fit within a single detector at the required resolving power of 2000. Thus, two channels are implemented in each spectrograph using a dichroic beamsplitter, with 3900-6100 Å reflected into the blue channel and 5900-9100 Å passed through into the red chan-

nel. The collimated beam diameter is 160 mm; larger would begin to require unreasonably large optics, while a smaller beam would require higher field angles into the cameras and therefore more difficult designs. With this beam diameter the cameras have a focal length of 240 mm, and reasonable grating designs can provide the necessary angular dispersion.

Figure 7 shows the optical layout of the SDSS spectrographs. Light enters each spectrograph through 320 fibers, which terminate at a curved slithead. The slithead positions the fiber ends on a radius concentric with the spherical collimating mirror, which operates at $f/4$ and produces a 160 mm diameter beam. The 45 degree dichroic beamsplitter reflects the blue portion of the bandpass ($\lambda < 6000\text{ Å}$) and transmits the red wavelengths ($\lambda > 6000\text{ Å}$). Located immediately behind the beamsplitter in each channel is a grism, consisting of a right angle prism with a transmissive surface-relief grating replicated on the hypotenuse. The dispersed light exits the grisms and enters all-refractive, eight-element, $f/1.5$ cameras. Each camera contains a single SITe $2\text{ k} \times 2\text{ k}$ CCD with $24\text{ }\mu\text{m}$ pixels. The camera demagnification from $f/4$ to $f/1.5$ produces fiber images that are just under 3 pixels in diameter. A fiber-center separation of 390 microns at the slithead produces six pixel center-to-center spectra spacing at the detector, sufficient to avoid crosstalk between spectra.

2.3.2. Slithead

Light enters the spectrograph through $180\text{ }\mu\text{m}$ diameter fibers, which terminate at the curved slit plate; see Section 2.2 for more details on the slithead design. The fibers are stacked vertically to form a tall, narrow, slit on a radius whose center of curvature coincides with that of the collimator. Additionally, the fibers are arranged in a fanlike pattern outward from the center of curvature toward the collimator, so that the central ray from each fiber strikes the collimator normal to the surface. Thus, the slit plate is at the focus of a one-dimensional Schmidt collimator.

2.3.3. Collimator

In combination with the curved slithead, the spherical collimator mirror forms a corrector-less Schmidt collimator. We found it possible to eliminate the corrector plate because our imaging requirements did not need the full Schmidt performance, and some of the reduced performance could be compensated for in the spectrograph cameras. Thus, our final collimator design is a single mirror with a spherical figure. The mirror itself is fabricated from a rectangular Hextek gas-fusion process borosilicate blank, 175 mm wide, 419 mm tall, and 73 mm thick. The planar blank is slumped by Hextek to near-net radius and then ground and polished to the final radius of 1264 mm. An enhanced silver coating was applied to the optical surface by Denton Vacuum. The collimator forms a pupil at the center of curvature of the mirror, and this is where the gratings are located in each channel in order to minimize their required size.

We examined alternatives to the Hextek blank including aluminum, eggcrate Zerodur, and monoliths. The gas fusion blank wins easily on weight and cost considerations. The compromise is in the thermal coefficient of




FIG. 7.— Optical layout of the SDSS spectrographs. Light enters each spectrograph through 320, 180 μm diameter, fibers, which terminate at a curved slit plate mounted inside the slithead. The slit plate positions the fiber ends on a radius concentric with the spherical collimating mirror, which operates at $f/4$ and produces a 160 mm diameter beam. The 45 degree dichroic beamsplitter reflects the blue portion of the bandpass ($\lambda < 6000 \text{ \AA}$) and transmits the red wavelengths ($\lambda > 6000 \text{ \AA}$). Immediately after the beamsplitter in each channel is a grism, consisting of a right angle prism with a transmissive surface-relief grating replicated on the hypotenuse. The dispersed light exits the grisms and enters all-refractive, eight-element, $f/1.5$ cameras. Each camera contains a single SiTe $2k \times 2k$ CCD with $24 \mu\text{m}$ pixels. The camera demagnification from $f/4$ to $f/1.5$ produces fiber images that are just under 3 pixels in diameter, resulting in 3 pixel tall spectra on the detector.

expansion, which is non-zero but much smaller than aluminum. An aluminum mirror would obviate the need to refocus when the temperature changes, but is expensive and the coatings are fragile. An eggcrate Zerodur mirror is immune to large temperature changes but is expensive and would require the most refocusing. A monolith would work well enough, but would not be significantly less expensive than the lighter and more thermally responsive gas fusion mirror.

To compensate for changes in instrument temperature with time, focus adjustment is provided by the collimator mount (see Section 2.4), which enables not only focus adjustment but tip/tilt adjustment as well. A pair of Hartmann doors located in front of the mirror (also described in Section 2.4) allow shifts in the collimator focus to be measured rapidly. Tip/tilt adjustment allows the collimator to be precisely coaligned to the cameras, or at least to an average of the two camera axes. Thus, the center fiber can be positioned at the center of the detector in the spatial direction, and the central wavelength can be positioned at the center of the detector in the spectral direction.

2.3.4. Beamsplitter

A dichroic beamsplitter divides the incident collimated beam, reflecting the blue portion of the bandpass ($\lambda < 6000 \text{ \AA}$) and transmitting red wavelengths ($\lambda > 6000 \text{ \AA}$). It is fabricated from BK7, is 271 mm wide, 229 mm tall, and 38 mm thick, with the dichroic coating applied to the incident surface. The coating reflects the blue light very efficiently ($R > 98\%$) and transmits the red light somewhat less efficiently ($T > 94\%$ average, including the reflection loss at the exit surface, which has a high performance broadband antireflection coating). The 10%/90% zone at the crossover wavelength is approximately 50 nm wide.

2.3.5. Gratings

The dispersing elements are grisms with zero angular deviation at 4960 \AA for the blue and 7400 \AA for the red. In our case, these are right angle BK7 prisms with a transmission grating replicated on the hypotenuse. While a reflection grating might have been used, the grism permits mounting the cameras close to the system pupil, which is about midway on the grating. With a reflection grating, the cameras have to be mounted away from the grating to avoid interference with the incoming beam, making them larger and more difficult to design. A plane transmission grating does not work because the diffracted angle is large, making geometric losses high (the groove facets are foreshortened) and forcing the blaze peak outside the optical band. Our configuration has little groove shadowing or foreshortening and results in high grating efficiency.

The ruling densities are 640 and 440 lines/mm for the blue and red grisms, respectively. Because master rulings of the size and groove angle needed did not exist, new masters were ruled by Hyperfine, Inc.

2.3.6. Cameras

The spectrograph cameras are $f/1.5$ with a 240 mm focal length and 16.5 deg field of view; see Figure 8 for a layout of the blue camera (the red camera layout is nearly identical). The all-refractive design was chosen to maximize throughput, because placing the detector or a secondary mirror in the unobstructed collimated beam, as would be the case for example with a Schmidt camera, would result in significant light loss. Both the red and blue cameras share the same optical prescription for the first four elements (i.e. the singlet and triplet). However, the doublets have unique prescriptions for the red

and blue designs, as to the field flatteners. In addition, given the uniqueness of each CCD surface curvature, the prescription for each field flattener lens pair is unique. See Table 1 for prescription details. Anti-reflective coatings are optimized for the respective bandpasses of each channel.

Harland Epps (UCSC) was contracted to design the cameras, using the LRIS spectrograph camera design as a starting point with some additional constraints. The most important was to limit aspheric surfaces to one per camera with only mild angles (LRIS had two steep aspheric surfaces). The reason for this was to reduce schedule risk and to increase the vendor pool. The camera employs eight lens elements arranged in five groups, including a contact triplet, a contact doublet, and three singlets. All the surfaces are spherical except for the air side surface of the second element in the doublet, which is a relatively mild asphere. Careful attention was paid to glass selection in order to maximize throughput at 390 nm, with five of the eight elements being either calcium fluoride (CaF2) or Ohara i-line glasses (which have > 98% internal transmission at 365 nm). Dow Corning Q2-3067 optical coupling grease was used to join the elements in the doublet and triplet. The resulting design, shown in Figure 8, is similar to LRIS but has a smaller diameter and larger field of view.




FIG. 8.— Optical layout of the SDSS spectrograph blue channel camera designed by Harland Epps. Diameters are in inches. The red channel camera layout shares the same eight-element configuration, though the prescription for the doublet and field flatteners differ for the two camera designs. Anti-reflective coatings are optimized for the respective bandpasses of each channel.

Once the fundamental camera design was in hand, it was tuned using Zemax to match surfaces with existing vendor test plates, optimizing spacings to melt data received from the glass vendor, and generating optical prescriptions for the two field flattening lenses to accommodate the different CCD curvatures. The problem of CCD curvatures was particularly annoying. The CCD surfaces were not flat to start and they changed their shape when they were cooled to operating temperature. Fortunately, they returned to the same shape every time they were cooled to operating temperature and most had a generally spherical shape that could be corrected with spherical optics. Unfortunately, they were all different so each CCD required a customized pair of field flattening lenses.

2.3.7. Optical Prescription

Table 1 shows the spectrograph optical prescription. Some of the surface descriptions can be deciphered from this example: “Doublet,first,back” refers to the doublet component, first lens element, back surface (closest to the CCD). The “Radius” is the radius of curvature in inches, negative implying a concave left surface. “Thickness” is the distance from the current surface to the next; positive to the right. All materials are from Ohara Glass except for CaF2 and the lens couplant, Dow-Corning Q2-3067.

2.3.8. Predicted Optical Performance

The following sections discuss the predicted optical performance of the SDSS spectrographs, including image quality, spectral resolution, and throughput. Measured performance is discussed in Section 4.

Image Quality— Spot diagrams for the red and blue channels are shown in Figure 9 and Figure 10, respectively. The spots are shown within a $67.5 \mu\text{m}$ diameter circle, representing the diameter of the imaged fiber on the detector. Each diagram covers the full respective bandpass of the channel, and field points cover the full length of the slit. The average RMS spot diameter for the red channel is 21.5 microns and the maximum RMS diameter is 29.2 microns. For the blue channel the average RMS diameter is 21.1 microns and the maximum RMS diameter is 27.3 microns. When convolved with the $67.5 \mu\text{m}$ fiber diameter, the geometric aberrations represented in these figures degrade the image, but to a minimal degree, and the resulting image width is well sampled by three pixels (72 microns) on the detector.

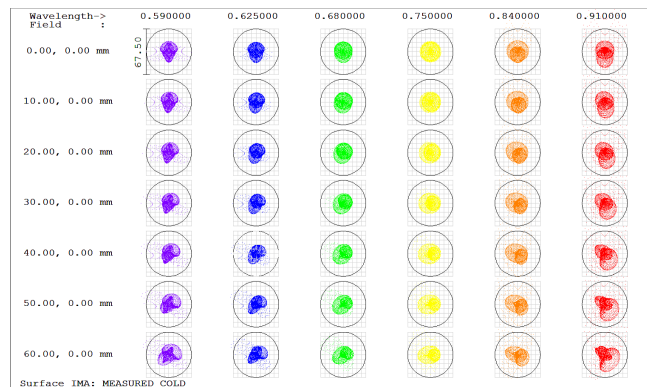


FIG. 9.— Spot diagram for the SDSS red channel. The average RMS spot diameter is 21.5 microns and the maximum RMS diameter is 29.2 microns.

Spectral Resolving Power— To analyze the spectral resolving power, defined as $\lambda/\delta\lambda$, where $\delta\lambda$ is taken to be the spectral FWHM of the fiber image on the detector in Å, the following procedure was used. In Zemax, a circular source the diameter of the fiber was placed at the center fiber location on the slit for the wavelength under consideration. Many thousands of rays were launched from this circular source with a uniform distribution, and the resulting image recorded on a simulated detector with pixels 1/4 the size of the $24 \mu\text{m}$ CCD pixels (in order to better sample the image). The simulated image data

TABLE 1
OPTICAL PRESCRIPTION FOR THE SDSS SPECTROGRAPHS

Surface	Radius (inches)		Thickness (inches)		Material
	Blue	Red	Blue	Red	
slithead	-25.2	-25.2	24.808	24.808	air
collimator mirror	-49.76	-49.76	-42.828	-42.828	air
prism	plano	plano	2.543	2.782	BK7
grating	plano(37°)	plano(39.5°)	2.070	2.300	air
singlet,front	-7.197	-7.197	-1.202	-1.202	CaF2
singlet,back	-71.400	-71.400	-0.100	-0.100	air
triplet,1st,front	-7.304	-7.304	-0.203	-0.203	LAL7
triplet,1st,back	-4.168	-4.168	-0.003	-0.003	Q2-3067
triplet,2nd,front	-4.168	-4.168	-2.501	-2.501	CaF2
triplet,2nd,back	8.903	8.903	-0.003	-0.003	Q2-3067
triplet,3rd,front	8.903	8.903	-0.201	-0.201	BSM2
triplet,3rd,back	-26.040	-26.040	-5.391	-5.352	air
doublet,1st,front	-4.350	-4.308	-0.200	-0.200	BAL35Y
doublet,1st,back	-3.436	-3.202	-0.003	-0.003	Q2-3067
doublet,2nd,front	-3.436	-3.202	-1.737	-1.850	FPL51Y
doublet,2nd,back (asph)	16.890	17.043	-2.107	-2.122	air
flattener,1st,front	5.372	5.337	-0.080	-0.080	PBM2Y
flattener,1st,back	29.999	23.951	-0.320	-0.270	air
flattener,2nd,front	4.832	4.505	-0.15	-0.15	LAL59
flattener,2nd,back	plano	plano	-0.151	-0.145	air
CCD	98.543	-75.592	0	0	silicon

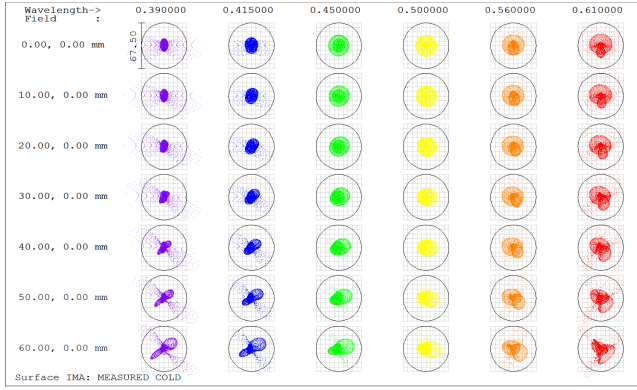


FIG. 10.— Spot diagram for the SDSS blue channel. The average RMS diameter is 21.1 microns and the maximum RMS diameter is 27.3 microns.

was analyzed to determine the FWHM, without collapsing the image. The FWHM thus determined was taken to be the resolution, and the resolving power is plotted in Figure 11.

In reality, the resolving power of the instrument is determined using collapsed spectra, so the prediction here is on the low side. However, the Zemax analysis assumes perfect optics and alignment. For this reason the predictions have been based on un-collapsed spectra, which is the conservative approach. The expectation is that the actual resolving power will be somewhat better than shown here.

Throughput— The total throughput, on-sky, for the SDSS spectrographs was predicted from an end-to-end component model as a function of wavelength. Included in the model were: atmospheric extinction, seeing (slit) losses, telescope, fibers, collimator, dichroic, grism, camera, and CCD. Atmospheric extinction was modeled at one airmass using a Palomar curve²⁷ scaled to Apache Point using mean values for photometric nights during 1998-2008 (Table 3 of Padmanabhan et al. 2008). Seeing

²⁷ www.sdss.org/dr3/instruments/imager/filters/index.html




FIG. 11.— Predicted resolving power, R , as a function of wavelength for the SDSS spectrographs. Predictions are derived from the FWHM of simulated, uncollapsed images computed using Zemax. In reality, the resolving power of the instrument is determined using collapsed spectra, so the prediction here is on the low side. However, the Zemax analysis assumes perfect optics and alignment. For this reason the predictions have been based on un-collapsed spectra, which is the conservative approach. The expectation is that the actual resolving power will be somewhat better than shown here.

losses were modeled using a double Gaussian PSF with a FWHM of 1'' centered in a 3'' aperture, resulting in a 5.5% throughput loss assumed flat across the bandpass. The telescope efficiency is based on measured data for CO₂-cleaned bare aluminum mirrors (Wilson, RTO II), along with simulated anti-reflection coating curves to match the specifications of the two wide field corrector lenses (a small overall effect). Measured curves were used for the dichroic, grism, camera coatings, and CCD quantum efficiency (QE). The manufacturer's curve for

Denton FSS99 silver was used for the collimator coating. Internal transmission curves for the camera glasses were obtained from the manufacturer’s data sheets. The fiber efficiency (85%) is based on lab measurements, representing an average value for the fibers measured and assumed flat across the bandpass. An additional 3% loss is included for focal ratio degradation beyond the f/4 beam the collimator was designed to accept. Figure 12 shows the individual component efficiencies used for this model along with the total expected throughput. Peak efficiencies in the blue and red channels are 17% and 22%, respectively.

2.4. Mechanical Design

2.4.1. Overview

The twin spectrographs mount to the SDSS telescope Cassegrain rotator adjacent to the focal plane with sufficient separation between the spectrographs to allow routine installation and removal of the imaging camera and fiber cartridges; see Figure 13. Compared to a bench-mounted configuration this approach greatly reduces fiber length, and the associated throughput loss. It also eliminates the potentially deleterious effects of repeated bending of the fibers. The obvious consequence of telescope-mounted spectrographs is, however, flexure induced motion of the spectra as the telescope tracks the sky, an issue that is mitigated by the use of a stiff, enclosed optical bench.

Each spectrograph has a mass of approximately 320 kg. The overall length, width, and height are 2210mm \times 1000mm \times 570mm, respectively. The optical layout of the instrument lends itself to a simple mechanical design, having only a few mechanisms and only one actively controlled optic, the collimator. To ease assembly and testing, the fiber slithead and all refractive components (i.e. camera lenses, the dichroic, and grisms) are mounted without adjustment in logical sub-assemblies, relying solely on machining tolerances for placement. Shims are located at strategic locations where opto-mechanical tolerances are tighter than could be achieved with reasonable machining practices. This basic approach has produced a robust, reliable instrument that has been straight-forward to maintain and upgrade. In the following sections we describe the mechanical design in more detail, with an emphasis on the opto-mechanical sub-assemblies.

2.4.2. Optical Bench

Figure 14 shows the spectrograph optical bench and the opto-mechanical subassemblies that interface to it. The optical bench is an enclosed aluminum (6061-T6) weldment with precision-machined interfaces to locate all five opto-mechanical subassemblies: the fiber slithead, the collimator, the central optics, and the red and blue channel cameras. The dual-channel optical layout naturally leads to the T-shaped configuration depicted in the figure. Both the red and blue cameras, as well as the collimator assembly, mount to exterior faces of the bench; one at the end of each leg of the T. The fiber slithead and the central optics assembly, which contains the dichroic and grisms, are mounted internally through access ports. At all of these opto-mechanical interfaces precision-machined reference datums and kinematic-locating features are employed to

ensure accurate, repeatable alignment of the optical system. A shim at the collimator interface provides a one-time coarse adjustment to compensate for the as-built collimator radius of curvature. Another shim where the slithead kinematic mount attaches to the bench provides a one-time adjustment to compensate for machining tolerances. Instrument control electronics and the CCD controller mounts to exterior walls of the bench. The guide camera, located on only one of the two spectrographs, mounts externally as well.

The principal considerations in designing the optical bench were ease of assembly, precise optical alignment, and the requirement to minimize flexure-induced motion of the spectra as the telescope tracks the field. With a tall closed-end, box-section design, internal and external gussets, and optimized placement of the mount locations to the rotator, flexure-induced motion of the spectra is kept to an acceptably low level during a typical exposure. Secondary considerations include the desire to reduce stray light from external sources and to keep exposed optics free of dust and debris.

2.4.3. Collimator Mount

Focus adjustment and registration of the spectra on the detector are accomplished by active adjustment of the collimator. The collimator is translated in piston to adjust focus in both channels, an adjustment that occurs routinely during operation to compensate for thermal drift, and flexure due to changes in telescope pointing and rotator angle. A pair of Hartmann doors located in front of the collimator permits shifts in the focus to be measured rapidly. Tip/tilt adjustment allows the collimator to be precisely coaligned to the cameras, or at least to an average of the two camera axes. Thus, the center fiber can be positioned at the center of the detector in the spatial direction, and the central wavelength can be positioned at the center of the detector in the spectral direction.

Details of the collimator mount are depicted in Figure 15. The mirror is supported from the back surface, at its center, by a circular membrane flexure, which constrains the mirror in-plane but allows out-of-plane compliance for tip, tilt, and piston adjustment. Motion is provided by three linear actuators (Physik Instrumente model #M-222.20), which provide sub-micron resolution, more than sufficient focus adjustment, adjustment of the spectra on the CCD, and flat-field dithers. Each actuator connects through a spring-pre-loaded universal joint to an Alloy42 steel pad bonded to the back of the mirror with 3M 2216 two-part epoxy. Three limit switch assemblies behind the mirror limit travel to ± 3 mm. The linear actuators, limit switches, and flexure are attached to a common mounting plate, which interfaces to the optical bench. Electrical connections from the actuators and limit switches are routed to a single connector on a custom printed circuit board mounted to the rear of the plate. A shim between the mounting plate and the bench facilitates one-time, gross focus adjustment. Accurate and repeatable assembly of the subsystem is made possible by two dowel pins (one round and one diamond shape) that engage hardened bushings pressed into the optical bench. Safe installation is facilitated by two handles for lifting, and four guide rods that flank the collimator and ensure the collimator clears the bench opening




FIG. 12.— Expected throughput for the SDSS spectrographs. The plot shows all contributors to the throughput model having a wavelength dependence. Not shown are those contributors with an essentially flat response across the bandpass: fiber transmission including Fresnel losses at the two faces (0.85), focal ratio degradation overfilling the collimator stop (0.97), and “slit” losses for 1'' FWHM seeing conditions modeled with a double Gaussian PSF (0.94). Not included in the model are losses due to centering and guiding errors. Overall system throughput, shown by the solid black curve, is expected to peak at about 17% in the blue channel and 22% in the red channel.




FIG. 13.— Rear view of the SDSS telescope depicting the location of the spectrographs. Here a fiber cartridge is shown retracted from the spectrographs. The twin spectrographs, each with a mass of 320 kg, mount to the back of the Cassegrain instrument rotator adjacent to the focal plane with sufficient separation between the spectrographs to allow routine installation and removal of the imaging camera and fiber cartridges.

as it is installed. A sheet metal cover surrounds the rear of the assembly, which protects vital components from physical and environmental damage, and reduces stray light. Finite element analysis was conducted early in the design phase of the project to ensure the loading condition described here would not distort the mirror to an appreciable degree. The results indicated that the distortion produced by gravity and the membrane load would be negligible.

2.4.4. Hartmann Doors

The Hartmann doors are a simple bi-fold design. Each door pivots about an upper and lower bushing pressed into the optical bench, and is driven by a 90° pneumatic rotary actuator (Bimba Pneu-Turn model PT-006090-A1M²⁸) located on the top surface of the bench. The doors are controlled by the spectrograph micro-controller and a bank of programmable solenoid valves (Clippard model number EMC-08²⁹) located in the spectrograph electronics box. Solid-state magnetic sensors (Bimba model number HSC-02) mounted to the actuators report the state of the doors (i.e. open or closed). Small manual flow control valves (Bimba model number FCP-1) at the inputs to each actuator set the speed of rotation.

2.4.5. Shutter

The spectrograph shutter is located just upstream of the dichroic on a dividing wall in the optical bench residing between the central optics and the slithead; see

²⁸ www.bimba.com

²⁹ www.clippard.com




FIG. 14.— Section views of the SDSS spectrographs. The T-shaped optical bench is an enclosed aluminum (6061-T6) weldment with precision-machined interfaces to locate all five opto-mechanical subassemblies: the fiber slithead, the collimator, the central optics, and the red and blue channel cameras. One of the two spectrographs also supports the guide camera. Electronics control chassis mount to the external walls of the bench. Three kinematic mounts on the top of the bench interface to the back of the Cassegrain instrument rotator.

Figure 14. This location is ideal since the entire band-pass can be blocked by a single, relatively small shutter. The shutter is a black anodized aluminum sliding door driven by a double-acting pneumatic cylinder (Bimba model BRM-02x-DXP). Slots in the Delrin door frame guide the door, making a light-tight seal when the door is closed. An oval window in the frame, slightly oversized to the beam, serves as a light baffle. Like the Hartmann doors, the shutter is controlled by the spectrograph micro-controller and solenoid valves on the Clippard EMC-08 board. Solid-state magnetic sensors (Bimba model number HSC-02) mounted to the cylinder indicate the state of the door (i.e. open or closed). Small manual flow control valves (Bimba model number FCP-1) at the inputs to cylinder set the speed at which the door opens and closes.

2.4.6. Central Optical Assembly

The beamsplitter and grisms are mounted in a single opto-mechanical structure, the central optics assembly. The assembly is kinematically mounted inside the optical bench on three gusseted posts integral to the weldment and located on the top surface of the optical bench, thus

the assembly is suspended from the top of the bench. The faces of the three posts are precisely machined establishing a planar reference that is square to the collimator mounting surface and the two camera mounting surfaces. Two locating sleeves centered about two of the three posts establish the in-plane location of the assembly. These machined references precisely locate the assembly relative to the slithead, collimator, and cameras. A large port in the sidewall of the optical bench provides access for machining the post surfaces and for installation of the assembly.

Figure 16 shows the details of the central optics assembly. The dichroic and both grisms are each located, without adjustment, by six machined reference surfaces. Spring plungers seat the elements against these reference surfaces and accommodate differential contraction between the glass optics and aluminum structure. A three-point-contact Alloy-39 block bonded to the top of each grism spreads the vertical load applied by a single large plunger embedded in the top plate, and provides a convenient lift point for installing the grism. To minimize tolerance stack-up and improve placement accuracy, all




FIG. 15.— Collimator mount for the SDSS spectrographs. The mirror is supported from the back surface, at its center, by a circular membrane flexure, which constrains the mirror in-plane but allows out-of-plane compliance for tip, tilt, and piston adjustment. Motion is provided by three linear actuators, which provide sub-micron resolution, more than sufficient for focus adjustment, adjustment of the spectra on the CCD, and flat-field dithers.

but one of the eighteen reference surfaces (six for each of three optics) is machined into a single component, the base plate, which interfaces to the optical bench. The single remaining reference surface, which controls the tip of the beamsplitter, is located on the top plate. Black anodized surfaces and thin light baffles at the exit faces of the assembly serve to mitigate stray light. The entire assembly has a mass of 39 kg.

2.4.7. Camera Opto-mechanics

The opto-mechanical design of the SDSS cameras, shown in Figure 17, was derived largely from the Norris spectrograph camera design (Cohen et al. 1988), a logical consequence of the two cameras having very similar optical designs, and the same optical designer, Harland Epps. Similar to the Norris camera, the SDSS camera prescription starts with a singlet, followed by a triplet, then a doublet, and finally two small field flattening lenses just in front of the detector. The lenses and lens-groups are mounted in athermal cells. External reference surfaces on each cell are machined true to the opto-mechanical refer-




FIG. 16.— Exploded view of the Central Optics Assembly. The dichroic and both grisms are each located, without adjustment, by six machined reference surfaces. Spring plungers seat the elements against these reference surfaces and accommodate differential contraction between the glass optics and aluminum structure. Kinematic mounting interfaces on the bottom of the base plate locate the assembly inside the optical bench.

ence surfaces that locate each lens, thus establishing lens concentricity from cell-to-cell and accurate placement of the lenses along the optical axis. The singlet and triplet cells are bolted together in series and mounted in an aluminum barrel along with the doublet cell. A steel inner barrel meters the distance between the triplet and doublet, the doublet being pressed against the inner barrel by a compression spring assembly behind the cell. A thick shim between the singlet and triplet allows a one-time adjustment to compensate for machining tolerance stackup. The field flatteners are integral to the cryostat, as described below.

The design and construction of the athermal cell design is described as follows. A metal ring with an appropriate coefficient of thermal expansion is bored oversize to the lens diameter. Six glass-filled Teflon plugs are lightly pressed into a hole pattern circumscribing the lens bore. Here the rear two elements are located by a single set of plugs given the small edge thickness of the middle, calcium fluoride, element. These plugs are then bored on a lathe to a diameter slightly oversized (0.05 mm on diameter) to the as-built lens diameter. In the same machining step the remaining critical surfaces are machined, thus achieving lens concentricity (within the limits of radial clearance), location along the optical axis and perpendicularity to the optical axis. The plug diameter is calculated such that the net change in the diameter of the finished bore is less than the diametrical clearance to the lens over the temperature range of interest, given the coefficients of expansion for the metal ring, the Teflon plug, and the glass lens. Where lens groups are packaged in a single cell, multiple sets of plugs are used. For the SDSS cameras, 6061-T6 aluminum was used for the singlet and doublet cells, and 330 stainless steel was used for the triplet cell. A Kapton shim rests between the lens locating face and the metal cell. An o-ring in the retaining ring contacts the first lens surface, providing

force to seat the lens and compliance to accommodate differential expansion. A rendered image of the triplet cell depicting these details is shown in Figure 18.

Camera focus is achieved by translating the dewar, which houses the detector and the two field flattening lenses, and attaches to a flange at the rear of the camera barrel. A shim between the dewar and flange provides a one-time coarse focus adjustment. Small focus adjustments to compensate for seasonal temperature variations are affected by turning the focus ring threaded onto the rear of the camera barrel. Turning the ring translates the dewar flange, which is pre-loaded against the ring by eight compression springs in the rear face of the camera barrel. Two guide pins, also in the rear face of the barrel, engage two linear bearings in the flange providing smooth axial motion with negligible lateral motion. The large diameter of the focus ring combined with a fine 32 thread-per-inch pitch allow focus resolution of roughly two microns. Marks engraved in one degree increments on the outer diameter of the ring provide a convenient method for metering adjustments.

The camera is bolted to the optical bench at the camera mounting flange, which is located at the camera center of gravity. Two pins (one dowel pin and one diamond pin) in the face of the flange engage two steel bushings in the bench, ensuring repeatable placement. Handles on the barrel and a separate, detachable guide-rail system, which bolts to the optical bench, facilitate safe removal and installation of the camera. The entire camera mass (including the dewar) is 40 kg.




FIG. 17.— Cross section of the SDSS blue camera highlighting the opto-mechanical design details. The opto-mechanical details of the red camera are identical. The design was derived largely from the Norris spectrograph camera design. The lenses and lens-groups are mounted in athermal cells. External reference surfaces on each cell are machined true to the opto-mechanical reference surfaces that locate each lens, thus establishing lens concentricity from cell-to-cell and accurate placement of the lenses along the optical axis. Manual focus adjustment is provided by a threaded focus ring at the back of the camera barrel. This adjustment is used to compensate for small focus drifts and to ensure the two channels are parfocal.




FIG. 18.— Exploded view of the SDSS blue camera triplet. The athermal cell design consists of a type 330 stainless steel ring with six glass-filled Teflon plugs lightly pressed into a hole pattern circumscribing the each lens bore. Here the rear two elements are located by a single set of plugs given the small edge thickness of the middle, calcium fluoride, element. These plugs are then bored on a lathe to a diameter slightly oversized (0.05 mm on diameter) to the as-built lens diameter. The net change in the diameter of the finished bore is less than the diametrical clearance to the lens over the temperature range of interest, given the coefficients of expansion for the metal ring, the Teflon plug, and the glass lens. An o-ring in the retaining ring contacts the first lens surface, providing force to seat the lens and compliance to accommodate differential expansion.

2.4.8. Dewars

The CCD is housed in a dewar and cooled using liquid nitrogen (LN₂); see Figure 19. The front half of the dewar contains the last two elements of the camera (i.e. the field flatteners), the detector, and the pre-amp board. O-rings locate the lenses on-axis and provide radial compliance for what is a small differential thermal contraction over the operating temperature range. Location of the CCD is adjustable in five degrees of freedom for centration, piston, tip and tilt. The detector columns are aligned to the slit by rotating the entire dewar relative to the back of the camera. The rear half of the dewar contains the LN₂ reservoir and is bolted to the front half of the dewar, a logical arrangement whereby the front half containing the optics, detector, and electronics can be assembled independent of the rear half containing the liquid nitrogen tank. A flexible copper braid provides the thermal connection between the reservoir and the detector. The space surrounding the reservoir and the remainder of the internal volume extending up to the rear of the last field flattener is evacuated. Detector temperature is controlled to 190 K using a small heater located on the cold-finger attached to the rear of the detector. The LN₂ reservoir is filled from the rear of the dewar through the center of a coaxial tube located on the reservoir axis and penetrating to the middle of the volume; this allows half of the 0.6 liter volume to be filled with liquid, regardless of gravity vector orientation. Gaseous boil-off vents through the outer tube.

The thermal design of the dewar is well optimized. The liquid nitrogen reservoir is fabricated from stainless steel and is well isolated from its surroundings. Aside from the detector cold strap, the only conductive paths to the reservoir are the thin-walled stainless vent tube and the G10 spider that provides radial support at the front end of the reservoir. Polished and gold-plated surfaces on the exterior of the reservoir and the radiation shield minimize radiative heat load. In operation the total dewar heat load is 4 Watts. An autofill system fills the dewar at roughly 3 hour intervals from a 10-liter storage dewar adjacent to the spectrograph.




FIG. 19.— Cross section of the CCD dewar. The liquid nitrogen dewar contains the last two elements of the camera (i.e. the field flatteners), the detector, and the pre-amp board. Location of the CCD is adjustable in five degrees of freedom for centration, piston, tip and tilt. In operation the total dewar heat load is 4 Watts. An autofill system fills the dewar at roughly 3 hour intervals from a 10-liter storage dewar adjacent to the spectrograph

2.4.9. Flexure

Minimizing flexure was a fundamental consideration throughout the design of the SDSS spectrographs. Mounted at the Cassegrain focus on the instrument derotator the gravity vector changes as the telescope tracks the sky. This variation in load causes gravity-induced deformations within the instrument to change with time and these changes lead to motion of the spectra on the detector. Any motion of the spectra can degrade resolution and increase cross-talk between the spectra. For this reason, the optical bench and the optomechanical assemblies (slithead, collimator, central optics assembly, and cameras) were designed to be extremely stiff.

The degree to which flexure matters is set by the cumulative exposure time plus the time it takes to perform a wavelength calibration. For SDSS, three fifteen minute exposures were taken per field. Considering the impact of flexure in both the spectral and spatial directions, and

allowing fifteen minutes for wavelength calibration, then leads to a stability criteria of 0.3 pixels in one hour, or 15° on the sky.

This overall budget was then divided equally between each of three subsystems: the slithead, the collimator, and the optical bench; each being allowed to contribute no more than 0.1 pixel shift in one hour. It was assumed, based on the details of the design, that the cameras and central optics assembly do not flex. It was also assumed, based on the designs put forth, that the slithead and collimator would not exceed the 0.1 pixel allowance. Hence a detailed analysis was not performed on these two subsystems. The collimator was analyzed for surface figure deformation as a function of gravity for the described mounting constraints, but these models assumed rigid actuators; something we knew was not true but was sufficient for the purpose of quantifying the change in surface figure.

The optical bench, however, did receive scrutiny. The primary concern being the deformation in the relatively thin plate structure with multiple heavy optical assemblies attached to it. To assess the stiffness of the optical bench a finite element study was conducted in collaboration with Swales Aerospace Corporation. The goal of the analysis was to quantify the shift in the spectra at the detector for a change in gravity vector orientation of 15° . We established a design goal a flexure of no more than 0.1 pixels for any 15° change.

The results showed that flexure in the bench is influenced largely by the placement of the three kinematic support points that attach the instrument to the rotator, as well as the thickness of the bench top plate. The optimal placement of the supports was determined to be at the edges of the bench; this set the final location for the mounts. In addition, the region of the bench around the mount points was stiffened by a thick plate, which became integral to the mount itself. The results indicated that with this configuration, and for rotations of 15 degrees about the red camera axis, the shift in the spectra would be much less than 0.1 pixels in both channels.

2.5. Detectors, Electronics and Data Acquisition

2.5.1. Detectors

The CCD detectors are thinned Tektronix/SITe SI424A 2048×2048 devices with square $24 \mu\text{m}$ pixels. The readout is done through two amplifiers at opposite ends of a single serial register. Readout noise is about 5 electrons and full well is about 150,000 electrons. A remarkable characteristic of these devices is a high quantum efficiency in the blue spectral region. The large pixels allow a modest demagnification ratio of 2.5, so the challenge is to put the full area of the detectors to good use. This requires a spectrograph camera with an exceptionally large field of view (16.5°) and the largest practical beam diameter (150 mm).

2.5.2. Detector Readout

The readout electronics are essentially identical to those used in the SDSS camera, which were described in detail in Gunn et al. (1998). In brief, the electronics consist of an in-vacuum preamp/clock driver using Burr-Brown OPA627 FET-input low-noise operational amplifier and a set of solid-state switches for DC restoration

and to generate the CCD clocks from supplied rail voltages and CMOS clock signals. This unit is physically and electronically identical to the ones used in the camera. The voltages and signals come to the dewar over a 37-wire cable, and the video is carried away over a 10-wire cable. These cables are routed to a *saddlebag* which is mounted to the dewar and carries three cards: 1) a bus receiver, which receives clock signals over an RS-485 bus serving both dewars on a single spectrograph and produces CMOS signals for the dewar and the other saddlebag cards; 2) a power distribution board, which regulates and filters incoming DC power for distribution to the saddlebags and dewar, and contains some of the temperature regulation circuitry; and 3) a complex 8-layer signal-chain/bias board which has circuitry on one side to generate all the needed rail voltages and on the other to implement CDS processing of the CCD video, digitize it, and drive an RS-485 output line to the controller. With the exception of some component values, these cards are identical to the ones used in the camera.

The RS-485 bus signals and the RS-485 digitized video to/from each dewar/saddlebag are carried to a single controller chassis for each spectrograph on two cables, the RS-485 on a standard 50-pin SCSI cable and the video/ADC signals on a ribbon cable terminated with DB-15s. The controller houses power supplies, using boards essentially identical to those used for the camera, the TDS Forth microprocessors used for all control functions and to generate the CCD clock waveforms, RS-485 driver/receiver circuitry, and circuitry to drive the FOXI fiber transmitters which send the (now multiplexed) digital data from the instrument to the data acquisition hardware some hundred meters away in the operations building. The camera uses two micros, one for executive control functions and one for waveform generation, but in the somewhat simpler spectrograph environment these functions are combined. Again, see the camera paper for details.

2.5.3. Data Acquisition

The data acquisition hardware for the spectrograph CCDs is a clone of the photometric camera system. The data are received on FOXI receivers, demultiplexed and stored on a ‘pool’ SCSI disk by a Motorola MVME167 single-board computer. Transfers from this VME backplane to the Linux host are handled by a custom interface designed and built at Fermilab.

On the Linux host, the observer’s program (SOS) actuates the system through remote procedure calls (RPC) over the Ethernet. The commands are few and simple, only those required to prep the CCD, expose, and read out the data. Once the data are off the CCDs and in the data system pool, the images are downloaded over the high speed VME link to the observer’s workstation where they are written to disk. At the end of the night, the SDSS data was archived and sent to Fermilab. A typical night of observing produces ~ 1.5 GB of data.

The drilling database is the list of objects on a plug-plate. It includes the name and drilling coordinates of each object and the field coordinates (α , δ , equinox) of the plug-plate. This information is needed at observing time so the telescope can be pointed and the plug-plate ID can be incorporated into the data file. This database is generated at Fermilab from the photomet-

ric imaging data and is delivered on the internet to the observer’s workstation at Apache Point well before the observations.

To keep track of which fiber went into which plug-plate hole, the plugging station is equipped with a device to map the plug-plate locations to the slithead after plugging is finished. The plugging database information is written directly to the observer’s workstation disk at plug time. This information is merged with the CCD data before it is written to disk.

2.5.4. Spectrograph Control

Spectrograph operations are controlled from a single program (called SOP) running on the observer’s Linux workstation. This program, the observer’s sole interface to the spectrographic system, is a clearinghouse for observing commands that translates observer’s requests into the commands required by the independent subsystems used for spectroscopy. These systems include: spectrograph microprocessor, telescope and guider, CCD data acquisition system, drilling database, and plugging station.

The spectrograph mechanical controls are handled by a Z-World Little Giant Z180 microprocessor board. This processor communicates with the observer’s workstation via RS-232A serial commands sent through an Ethernet terminal server located at the telescope. This processor handles all motion control (shutters, Hartmann masks, collimator tip, tilt, and focus) and also reports ambient and instrument temperatures for focus compensation.

The observer’s software communicates fully with the telescope control system via an Ethernet telnet connection. When a new plug-plate is locked on the telescope, the spectrograph relays the plug-plate ID to the observer’s software, which looks up the coordinates in the plug-plate database. The observer’s software commands the telescope to move to the mid HA position for flat field and wavelength calibration, then to the field for precise positioning, scaling, and focusing. When these operations are finished, the observer’s software starts the guider, opens the shutters, and begins observing.

2.6. Calibration

Calibration for the spectrographs, both flat-fielding to tie the responses of the fibers together for flux calibration, and the use of spectral lamps to provide wavelength calibration, are provided as part of the telescope structure. The screens for these functions as well as the lamps are mounted on the windbaffle structure, which serves as a dome/enclosure for the SDSS telescope (see e.g., Gunn et al. 2006).

At the top of the windbaffle structure is a mechanism incorporating eight lightweight aluminum-honeycomb sector-shaped panels mounted on shafts and driven by DC gearmotors. These panels can be closed to cover the annular aperture of the telescope or opened to allow observing. The bottoms of these panels are coated with a special wide-spectrum Lambertian paint (PolarKote), providing a highly uniform diffuse reflective surface.

The panels are illuminated in quarters by a bank of three projectors in each corner of the windbaffle structure just above the central structure which supports the primary mirror and the altitude bearings. The projectors project a quarter-annulus which fills the pupil of the

telescope and avoids illuminating the secondary structure directly. One of these units in each of the three-projector assemblies is illuminated by a 3100K quartz-iodine lamp and is used for flat fields; another houses a mercury-cadmium lamp, which has a moderately rich emission spectrum in the blue and green, and the last houses a neon lamp, which is used primarily to calibrate the red region of the spectrum. This system is also used in BOSS, the only change being the replacement of the neon lamp by a carefully constructed neon-argon tube to provide argon lines farther to the red than the neon spectrum reaches.

The flat-field lamp is fitted with a combination of glass Schott filters (2mm FG3 + 2mm BG14) in an attempt to mitigate the very rapid falloff of the lamp spectrum to the blue, so that reasonable flat fields can be obtained with good S/N in one exposure with no danger of saturation. This approach was moderately successful, but at the expense of producing a rather ‘lumpy’ spectrum.

Since the spectrographs are mounted on the telescope, they flex, and so it is necessary to be careful about flat fields. If the image of a fiber shifts one pixel during an exposure, the registration of the object spectrum and the flat field used to calibrate it is quite poor, and poor reductions will result. For this reason, special flat fields are recorded in which the collimator mirror is dithered slightly in angle along the spatial direction so as to produce wider traces than the normal fiber ones to allow pixel-to-pixel variations to be dealt with properly even if the spectrum moves slightly. This adjustment is accomplished by the tip/tilt/focus motors on the collimator, which have more than adequate resolution for this task.

2.7. Guiding

Guiding for the SDSS spectrograph is accomplished by the use of coherent imaging fiber bundles which are connected to the plugplates at the locations of guide stars using ferrules which are externally identical to the data fiber ferrules. The imaging bundles used in SDSS-I and -II were Sumitomo fibers of two sizes which projected to 7" and 11" in the focal plane. There were nine 7" bundles and two 11" bundles. The larger ones were primarily intended for acquisition of the guide stars, although they were also used for normal guiding as well after all the stars were centered.

The plugging ferrules were modified slightly for this purpose; it is necessary, clearly, to orient the guide fibers rotationally so the guider will know what to do when a star is off-center. This is accomplished by mounting a collar with a small pin on the exterior diameter of the ferrule. A hole, smaller than the ferrule holes, receiving the pin is drilled into the plate at the time it is drilled, and the plate pluggers rotate the guide ferrule to drop the pin into the hole. The guide fibers are *not* randomly plugged, but assigned to marked holes in the plate.

The guide fibers all are routed to eleven holes in a small block into which they are cemented and polished. That block is imaged through a back-to-back pair of 50 mm f/1.2 Nikon camera lenses with a 2 mm Schott BG38 filter in the parallel beam between them onto the CCD in a Roper Scientific (Photometrics) Sensys camera. The CCD is a Kodak KAF0401E, which has an 768×512 array of $9 \mu\text{m}$ pixels. The detector was normally run in bin-by-2 mode, which resulted in $18 \mu\text{m}$ ($0.3''$) pixels.

The camera has a thermoelectric cooler, which limits the CCD temperature to 10°C , but is otherwise not temperature controlled. The camera came equipped with a SCSI interface, and was connected to an Apple Macintosh computer with a commercial SCSI-to-fiber-to-SCSI interface.

The images are buffered on the Mac and sent by fiber ethernet to the control room where they are received and processed by the Linux SOP computer. A PID servo was implemented in that machine to generate correction signals to the telescope control computer. Display code to allow the observers to monitor the performance of the guider, either by generating basically a raw image of the fiber block or a display in which the images of the fibers are placed approximately where they are on the sky and oriented properly, was written.

The performance of the system was never completely satisfactory; the CCD was (essentially) uncooled, was noisy, had many dark defects, and its mounting was not mechanically satisfactorily stable. Halfway through the survey new code was written that handled centroiding properly even if the star was at the edge of the fiber image, and explicitly used dark frames and flat fields to correct the raw images. This enhancement helped, but a new and better camera was clearly needed, which did not arrive until the beginning of the BOSS survey.

3. BOSS UPGRADE DESIGN

3.1. BOSS Requirements

The SDSS camera and spectrographs were operated for five years from 2000 - 2005, followed by a three-year extension known as SDSS-II. In 2005, the Astrophysical Research Consortium, which owns and operates the Apache Point Observatory, put out a call for proposals to operate the 2.5m Sloan telescope and its instruments after 2008. In response to this announcement, the Baryon Oscillation Spectroscopic Survey (BOSS; Schlegel et al. 2009; Dawson et al. 2012) was proposed and awarded five years of dark observing time beginning in 2009. SEGUE-II, a one-year extension of the Sloan Extension for Extra-Galactic Understanding and Exploration (SEGUE; Yanny et al. 2009) was approved for one observing season to use the SDSS spectrograph from mid-2008 to mid-2009. The primary scientific objective proposed by the BOSS project was to map the baryon acoustic peak out to redshift $z = 0.6$ with percent-level precision. In addition, BOSS would pioneer a new technique to detect large scale structure in the spectra of distant quasars by observing the so-called Lyman- α forest, a series of absorption features caused by intervening clouds of hydrogen gas.

To achieve the BOSS scientific goals, a major upgrade of the SDSS spectrographs was required. In particular, BOSS needed a higher multiplex factor in the number of fibers per exposure, and higher optical throughput. To achieve this, the BOSS proposal included plans to build a new fiber system and to upgrade the optical system and cameras of the original SDSS spectrographs.

BOSS replaced the SDSS spectrograph cartridges, with 640 fibers of 3" diameter (180 micron), with cartridges having 1000 fibers with 2" diameter (120 micron). Smaller fibers were necessary in order to increase the number of spectra per exposure while re-using most

of the spectrograph optics, and the the smaller fibers would also result in better S/N measurement for high redshift objects, by reducing the contribution from sky background.

To improve throughput, the SDSS gratings were replaced with state-of-the-art large format volume-phase holographic (VPH) gratings, and the CCDs were upgraded with modern devices with higher quantum efficiency and smaller pixels to match the smaller fiber diameter. On the blue side of the spectrograph, new e2v³⁰ CCDs were installed for improved quantum efficiency in the far blue while fully depleted devices from Lawrence Berkeley National Laboratory (LBNL) were installed for improved quantum efficiency near 10,000 Å (Holland et al. 2006). The BOSS forecasts for new cosmological constraints assumed an improvement in throughput that was a function of wavelength, with a peak improvement of about a factor of two. The improved throughput was necessary in order to complete the survey in five years, while obtaining spectra with sufficient S/N for high efficiency (> 95%) redshift identification of faint galaxies.

Another requirement for BOSS was a modest increase in the wavelength range of the existing spectrographs. The Lyman- α transition is at 1216 Å in the quasar rest frame, and therefore the Lyman- α forest of a $z = 2.2$ quasar extends blueward from about 3900 Å. The minimum wavelength requirement for BOSS was set at 3700 Å, slightly lower than the original SDSS cutoff at 3900 Å, so as to detect at least part of the forest for quasars at redshift $z = 2.2$. The red cutoff was also extended, from ≈ 9000 Å to $\approx 10,000$ Å, so as to reliably detect the 4000 Å break feature in galaxies up to a maximum redshift of $z = 0.8$. The blue and red extensions in wavelength coverage were enabled by improvements in CCD technology.

3.2. Fiber System Upgrade

A major part of the SDSS spectrograph upgrade for BOSS was the acquisition of nine new fiber cartridges with 1000 fibers to replace the original SDSS cartridges which had 640 fibers. For BOSS, the cartridge design, including the castings themselves, is identical to that used for SDSS with the exception of revised slitheads to accommodate more fibers having a smaller diameter. The fiber mapper system follows the same philosophy as that used for SDSS, and fiber performance was tested in the same manner as the original SDSS fibers. For more detail on the cartridge design, the fiber mapper, and the method for assessing fiber performance see Section 2.2. Here we describe the BOSS fiber system upgrade, which for the most part is limited to the implementation of a larger number of smaller fibers.

3.2.1. Fibers

The optical fiber selected for BOSS is a step-index, all-silica, UV-enhanced fiber. It is similar to the fiber originally chosen for SDSS, except that it has a smaller core diameter of 120 microns. The core is surrounded by 25 microns of a lower-index cladding with a cladding-to-core ratio of at least 1.4, which confines the radiation by total internal reflection. The cladding is protected by a

polyimide sheath, or buffer, bringing the outer diameter to 190 microns. The tolerance on the outer diameter was specified at ± 3 micron, with a maximum decentering of the core relative to the outer diameter of 6 microns. These tolerances were tighter than the ± 5 micron outer diameter that was specified for BOSS prototype fibers with a larger outer diameter of 197 micron. The fiber was manufactured by Polymicro Technologies Inc. and is known as FBP 120-170-190. The total length of each fiber is 1830 ± 25 mm.

The design of the fiber harnesses for BOSS follows that of the original SDSS fiber harnesses as described in Section 2.2, with the same overall dimensions and components. The same manufacturer, C-Technologies, Inc, was selected during a competitive bidding process. As in the original SDSS design, a fiber harness consists of 20 fibers, with one end mass-terminated by gluing the fibers into a metal v-groove block, while the free ends of the fibers are individually terminated in stainless steel ferrules. The v-groove blocks align the fibers on 240 μm centers; the v-groove blocks are then glued to the slit heads. The v-groove blocks were machined at the University of Washington. The stainless steel ferrules were manufactured by Swiss Screw Products, Inc. After the fiber was inserted into the ferrule and glued using a low-shrinkage adhesive, the end of the fiber was polished, removing a small amount of material and maintaining strict tolerances on planarity. The ferrule end was designed for insertion into the holes drilled in the plug plates, where it is held by friction. A jacket of nylon tubing on the ferrule end protects the fibers from breakage during repeated plugging and unplugging operations.

3.2.2. Fiber Performance

To measure the FRD characteristics of BOSS fibers as configured for science observations (as well as to characterize the effective exit pupil of the telescope-spectrograph system), the testing setup diagrammed in Figure 20 was constructed. This system was used to illuminate 18 fibers one at a time within a spare BOSS harness with a white-light input beam matched to the $f/5$ SDSS 2.5m telescope beam, *including the 27% areal obscuration that removes the central rays from the cone of light incident on the telescope focal plane*. The ferrule of each fiber under investigation was held in place using an aluminum plate drilled with a decommissioned SDSS-III mill bit. A tip-tilt bracket was used to adjust the fibers for telecentricity with respect to the input beam. The diverging output beam was imaged using an SBIG ST-8300M CCD camera with 3326×2504 pixels of $5.4 \mu\text{m}$ pitch, with the detector plane located a distance of 4.7 ± 0.1 cm from the output termination of the fibers in the v-groove block. The resulting beam-profile images show a characteristic “fuzzy donut” of the telescope pupil convolved with the FRD within the fiber. The beam patterns were flux-centroided in the images, and the fractional flux enclosed as a function of distance from beam center was computed for each of the 18 fibers. The resulting curves are shown in Figure 21, with respect to an abscissa of beam half-angle in degrees. The mean curve for all 18 fibers is shown in thick black. We confirm that with an $f/4$ spectrograph collimator, essentially all of the emergent flux is captured by the optical system.

³⁰ <http://www.e2v.com/>




FIG. 20.— BOSS fiber focal-ratio degradation test setup diagram: (1) White LED array; (2) Diffuser tube; (3) Diffusing screens; (4) Light shroud; (5) Newtonian reflector telescope stopped to $f/5$ with 27% central areal obscuration; (6) Fiber input bracket; (7) BOSS optical fiber; (8) V-groove block with fiber termination; (9) SBIG ST-8300M CCD camera; (10) Data-acquisition computer; (11) Optical breadboard (24×48 in). System shown approximately to scale.




FIG. 21.— Focal-ratio degradation curves for 18 tested BOSS fibers (thin gray), with average curve (thick black).

3.2.3. Slithead

The BOSS slitheads are nearly identical in design to the original SDSS slitheads. They consist of a thin, stiff steel plate with a curved edge of radius 640 mm. There are two slitheads per cartridge, each having 500 fibers. The 120 μm diameter fibers are mounted in groups of twenty in v-groove blocks, with 25 v-groove blocks being glued to each slithead. The center-to-center spacing between fibers on adjacent v-groove blocks is 624 microns, compared to 260 microns between fibers within a v-groove block. The total length of the arc is 138.6 mm from outside edge to outside edge of the first and last fibers, slightly taller than the original SDSS design.

3.3. Optical Upgrades

3.3.1. Optical Design Overview

The optical design of the BOSS spectrograph is nearly identical to the original SDSS design. For BOSS the bandpass is extended in both channels to cover $3600 < \lambda < 10000 \text{ \AA}$. Smaller, 120 μm fibers replace the original 180 μm fibers used in SDSS. The smaller fibers are a better match to BOSS targets and allow more fibers

per spectrograph. Additionally the replicated surface relief gratings used in the SDSS design were replaced with volume phase holographic (VPH) gratings, and modern CCDs with higher quantum efficiency and smaller pixels were installed; the combined effect being an impressive 70% peak instrumental efficiency, a significant improvement over the 45% peak of the SDSS design.

Figure 22 shows the BOSS optical layout. Light enters the spectrograph through 120 μm fibers, which terminate at a curved slithead in the same manner as the original SDSS slithead. The collimator, unchanged from the original design, collimates the $f/4$ beam producing a 160 mm diameter beam, reflecting it back toward the dichroic beamsplitter. Light blueward of 640 nm is reflected into the blue channel, with the remaining light being transmitted into the red channel. A VPH grism in each channel disperses the light, which is then imaged by upgraded cameras, each containing a single $4k \times 4k$ CCD with 15 μm pixels. The use of smaller fibers and correspondingly smaller pixels, along with the same collimator/camera demagnification preserves the scale of the spectra on the detector, thus yielding 3 pixel wide spectral profiles separated by 6 pixels center-to-center; identical to the SDSS configuration.

3.3.2. Collimator

The BOSS survey started in 2009 with the same collimators used for SDSS. At that time protected silver coatings offering significantly higher reflectivity at the blue end of the spectrum had been realized on small optics. Given the obvious benefit of such a coating, the BOSS project set out to re-coat the collimators, beginning with the spare collimator originally produced for SDSS. However, realization of such a coating on a large surface such as the collimator was challenging due to the process control required to achieve very uniform layer depositions over a large surface area. Attempts by one vendor ultimately failed after more than a year of trial and error. A successful coating was achieved by Infinite Optics in late 2011, and newly coated collimators have been installed in both spectrographs.

3.3.3. Beamsplitter

The dichroic beamsplitter in each spectrograph was replaced for BOSS given the expanded bandpass of the instrument. For BOSS, wavelengths shorter than 605 nm are reflected by the new coating, a 15 nm shift toward the red as compared to SDSS. Although the BOSS bandpass is extended much more into the red than the blue, the crossover wavelength was shifted only a small amount in order to keep the transition region between channels in a scientifically less interesting region. In the process the entire beamsplitter, i.e., substrate and coating, were replaced; a logical consequence of the desire to keep the instrument operational while the new coatings were applied, and to mitigate the risk of potentially taking the instrument offline for an extended period should stripping the original coating lead to damage of the substrate.

An additional, and significant, consideration in replacing the beamsplitters was the desire to obtain a denser coating using ion-assisted deposition, which would be more immune to wavelength shifts with variations in humidity. In particular, the transition region of the SDSS beamsplitters shifted with changes in humidity inside the




FIG. 22.— Optical layout of the BOSS spectrographs. Light enters each spectrograph through 500, 120 μm diameter, fibers, which terminate at a curved slit plate mounted inside the slithead. The slit plate, which is integral to the slithead, positions the fiber ends on a radius concentric with the spherical collimating mirror, which operates at $f/4$ and produces a 160 mm diameter beam. The 45 degree dichroic beamsplitter reflects the blue portion of the bandpass ($\lambda < 6050 \text{ \AA}$) and transmits the red wavelengths ($\lambda > 6050 \text{ \AA}$). Immediately after the beamsplitter in each channel is a VPH grism, consisting of a VPH grating sandwiched between a pair of prisms (the grating is tilted 3.5° to steer the Littrow ghost off the detector). The dispersed light exits the grisms and enters all-refractive, eight-element, $f/1.5$ cameras. Each camera contains a $4\text{k} \times 4\text{k}$ CCD with $15 \mu\text{m}$ pixels. The camera demagnification from $f/4$ to $f/1.5$ produces fiber images that are just 3 pixels in diameter, resulting in 3 pixel tall spectra on the detector.

instrument, which occurs due to the periodic insertion of new cartridges. Each time a cartridge is installed the slithead door is opened, allowing moist air to displace the dry air injected into the bench to protect the calcium fluoride singlet in the camera. Once the door closes the humidity level slowly reduces back to its original near-zero level, and the dichroic transition shifted during this entire humidity excursion.

3.3.4. Gratings

As in the SDSS spectrographs, a grism immediately follows the beamsplitter in each channel. However, the significant increase in efficiency offered by VPH gratings convinced us to abandon the SDSS replicated surface relief grism design in favor of a VPH grism. The grism consists of a VPH transmission grating sandwiched between two prisms; the use of prisms is necessary to preserve the straight-through beam layout of the SDSS design.

The BOSS design utilizes a ruling density of 400 $1/\text{mm}$ in the red channel and 520 $1/\text{mm}$ in the blue channel. The gratings were designed independently at Johns Hopkins University (using the Zemax raytrace program³¹) and at Kaiser Optical Systems, Inc. (KOSI), who built the gratings. Both the VPH grating substrates and the prisms are made from BK7 or equivalent glass. The blue channel components made from Ohara Corporation's BSL7Y glass, a BK7 equivalent with enhanced UV transmission. The prism apex angles (total for the two-prism assembly) are 28° and 35.6° for the blue and red gratings, respectively.

A phenomenon now well-known to builders of VPH grating-based spectrographs, is that of the Littrow, or recombination ghost (i.e., Littrow ghosts; Burgh et al. 2007). This is an undispersed (white-light) ghost image

of the entrance slit or fibers of the instrument, and appears in the image plane at the location of the blaze or Bragg wavelength. Because the ghost is undispersed, its intensity can be quite bright relative to the nearby features of interest in the spectrum. This particular type of ghost is present in any spectrograph operating in the Littrow configuration ($\alpha = \beta$) but plagues in particular instruments with VPH gratings as they are so often used at or near the Littrow mode in order to obtain the highest efficiency possible.

To better understand the implications of the recombination ghost for the BOSS spectrographs, we undertook a rather extensive raytrace analysis of the possible ghost paths initiated by a reflection off the detector. In addition, we analyzed the paths of the zero-order light and the undesired diffracted orders with any appreciable efficiency predicted by KOSI's rigorous coupled-wave analysis of the grating designs. In order to move the recombination ghost away from the useful part of the spectrum, the grating must be moved away from the Littrow configuration. One way to do this with a VPH grating is by rotating the fringes within the gelatin layer during exposure, to be off-normal with respect to the plane of the grating itself. The grating is rotated by the same angle in use such that the fringe orientation is returned to the Bragg condition where efficiency is highest, but the grating is no longer at Littrow. Our analysis indicated a rotation of three degrees would be sufficient to move the ghost completely off the detector, and we settled on 3.5° degrees to provide a bit of margin. To accommodate the tilted grating each prism deviates from a right angle by the same 3.5° degree tilt, with one prism having an apex angle that is 3.5° degrees larger and the other 3.5° degrees smaller than the right angle design of an untilted grating. We chose this approach, rather than using

³¹ www.zemax.com

right angle prisms and rotating the entire grism assembly as a whole, due to extreme space constraints within the upgraded central optics assembly. Figure 23 shows the nominal, first order light path in the upper panel, and in the lower panel the recombination ghost path arising within the grating layer.

For BOSS, the VPH red grating peak efficiency is 82% as compared to 64% for SDSS, and the blue grating peak efficiency is 80% as compared to 50% for SDSS.

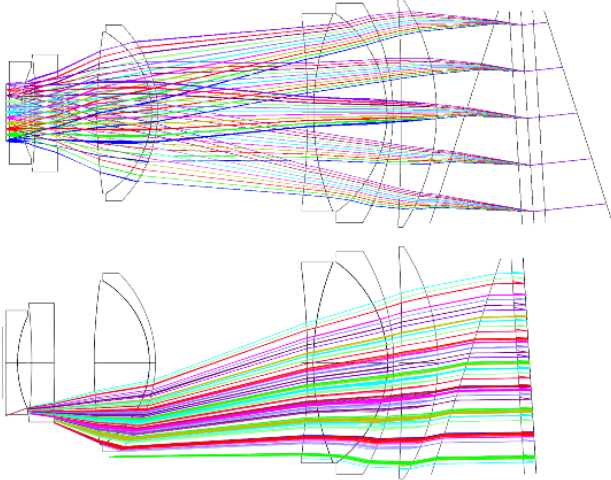


FIG. 23.— Raytrace layouts illustrating the light path for the recombination (Littrow) ghost. In these layouts the CCDs are on the left, the VPH grism assembly (or part thereof) is on the right. The desired, first order diffracted light (top) travels right to left through the grism and camera lenses to the CCD. The recombination ghost arising within the grating layer (bottom) is initiated by a reflection off the CCD, which sends some small percentage of the light back through the camera in a left-to-right direction. The camera recollimates this light, which then encounters the grating for a second time. Some of this light is reflected and recombined by the grating and travels back through the camera to the CCD. It is this last, right-to-left pass that is shown in the layout, the other passes are omitted for clarity. Note how the grating is rotated by 3.5 degrees with the prism assembly in order to direct the recombination ghost entirely off the CCD.

3.3.5. Cameras

Both red and blue camera optical designs were reoptimized for BOSS. The revised designs use the original SDSS prescriptions for the first six elements with the same multiplet configuration: a singlet, followed by a contact triplet, and then a contact doublet. As in the original design this sequence of elements is followed by two field flatteners, which reside just in front of the detector. For BOSS the spacings between these groupings were modified slightly, and the field flatteners were completely redesigned for improved imaging performance and larger focal plane arrays. These changes were motivated in part by the expanded wavelength range in each channel, and in part by new, larger ($4k \times 4k$ as compared to $2k \times 2k$), detectors having smaller pixels ($15 \mu m$ for BOSS as compared to $24 \mu m$ for the SDSS design) and working over a slightly larger field. Hence imaging performance requirements for BOSS are more stringent than for SDSS. Nonetheless the BOSS design uses the same

basic eight-element configuration as the original SDSS cameras.

3.3.6. Optical Prescription

Table 2 shows the optical prescription for the BOSS spectrograph. The surface descriptions are the same as in Table 1. As with the SDSS spectrographs, all materials are from Ohara Glass except for CaF2 and the lens couplant, Dow-Corning Q2-3067.

3.3.7. Predicted Optical Performance

The following sections discuss the predicted optical performance of the BOSS spectrographs, including image quality, spectral resolution, and throughput. Measured performance is discussed in Section 4.

Image Quality— Spot diagrams for the red and blue channels are shown in Figure 24 and Figure 25, respectively. The spots are shown within a $45 \mu m$ diameter circle, representing the diameter ($2''$) of the imaged BOSS fiber on the detector. Each diagram covers the full respective bandpass of the channel, and field points cover the full length of the slithead. The average RMS spot diameter for the red channel is 18.3 microns and the maximum RMS diameter is 27.5 microns. For the blue channel the average RMS diameter is 19.6 microns and the maximum RMS diameter is 30.9 microns. Compared to the SDSS image quality, the BOSS design is slightly better on average, and this is over an expanded bandpass and larger field of view (taller slit). However, the fiber diameter is smaller for BOSS as are the detector pixels. Nonetheless, when the geometric aberrations are convolved with the fiber diameter the resulting FWHM is still very near 45 microns, which matches three rows of pixels on the detector.

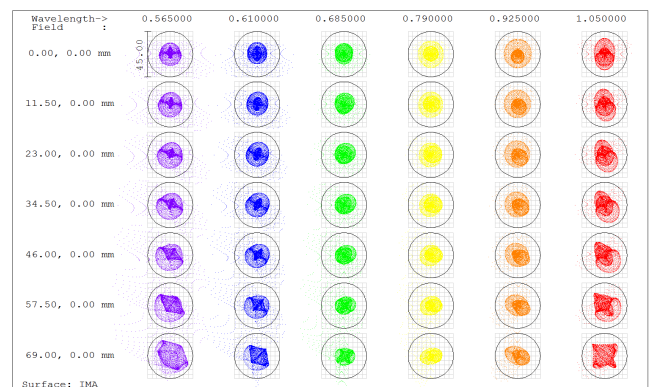


FIG. 24.— Spot diagram for the BOSS red channel. The average RMS spot diameter is 18.3 microns and the maximum RMS diameter is 27.5 microns

Spectral Resolving Power— Spectral resolving power predictions for the BOSS spectrograph are shown in Figure 26. These predictions were derived in the same manner as for the SDSS spectrograph resolving power discussed in Section 2.3. Comparing the BOSS and SDSS predictions, the blue channel is slightly lower than the SDSS design, while in the red channel the resolving power is slightly higher. The improved resolving power in the red

TABLE 2
OPTICAL PRESCRIPTION FOR THE BOSS SPECTROGRAPHS

Surface	Radius (inches)		Thickness (inches)		Material	
	Blue	Red	Blue	Red	Blue	Red
slithead	-25.2	-25.2	24.808	24.808	air	air
collimator mirror	-49.76	-49.76	-42.828	-42.828	air	air
prism	plano(-14.0°)	plano(-17.8°)	-1.323	-1.473	BSL7Y	BK7
cover plate	plano(-3.5°)	plano(-3.5°)	-0.315	-0.315	BSL7Y	BK7
grating	plano(-3.5°)	plano(-3.5°)	-0.315	-0.315	BSL7Y	BK7
prism	plano(14.0°)	plano(17.8°)	-1.314	-1.575	BSL7Y	BK7
singlet, front	-7.197	-7.197	-1.202	-1.202	CaF2	CaF2
singlet, back	-71.400	-71.400	-0.260	-0.274	air	air
triplet, 1st, front	-7.304	-7.304	-0.203	-0.203	LAL7	LAL7
triplet, 1st, back	-4.168	-4.168	-0.003	-0.003	Q2-3067	Q2-3067
triplet, 2nd, front	-4.168	-4.168	-2.501	-2.501	CaF2	CaF2
triplet, 2nd, back	8.903	8.903	-0.003	-0.003	Q2-3067	Q2-3067
triplet, 3rd, front	8.903	8.903	-0.201	-0.201	BSM2	BSM2
triplet, 3rd, back	-26.040	-26.040	-5.187	-5.107	air	air
doublet, 1st, front	-4.350	-4.308	-0.200	-0.200	BAL35Y	BAL35Y
doublet, 1st, back	-3.436	-3.202	-0.003	-0.003	Q2-3067	Q2-3067
doublet, 2nd, front	-3.436	-3.202	-1.737	-1.850	FPL51Y	FPL51Y
doublet, 2nd, back	16.890(asph)	17.043(asph)	-1.736	-1.371	air	air
flattener, 1st, front	11.603	130.376(asph)	-0.589	-0.740	BSM51Y	CaF2
flattener, 1st, back	plano	-19.821	-0.370	-0.478	air	air
flattener, 2nd, front	3.905	3.739	-0.240	-0.401	BSM51Y	S-LAH59
flattener, 2nd, back	plano	plano	-0.125	-0.120	air	air
CCD	plano	plano	0	0	silicon	silicon

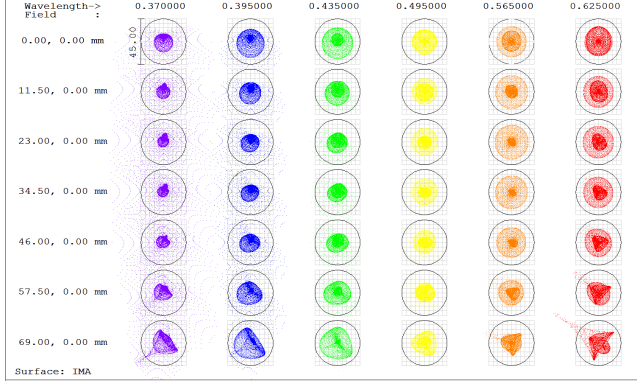


FIG. 25.— Spot diagram for the BOSS blue channel. The average RMS diameter is 19.6 microns and the maximum RMS diameter is 30.9 microns.

channel is the result of a larger detector with more pixels, more than offsetting the expanded wavelength coverage. In the blue channel the grating dispersion was deliberately chosen to underfill the detector in the spectral direction, thus maximizing the signal-to-noise ratio while providing sufficient resolving power to carry out the desired science.

We also present a two-dimensional contour image (Figure 27 that represents the RMS linewidths a function of detector position for the BOSS CCDs. Note that the linewidths are moderately degraded at the shortest wavelengths of the blue channel near the ends of slit.

Throughput— As was the case for the SDSS spectrographs, the total throughput, on-sky, for the BOSS spectrographs was predicted from an end-to-end component model as a function of wavelength. Included in the model were: atmospheric extinction, seeing (slit) losses, telescope, fibers, collimator, dichroic, grism, camera, and CCD quantum efficiency (QE). Atmospheric extinction was modeled at one airmass as before. Seeing losses were modeled using a double Gaussian PSF with a FWHM




FIG. 26.— Predicted resolving power, R , as a function of wavelength for the BOSS spectrographs. As with the SDSS design, predictions are derived from the FWHM of simulated, uncollapsed images computed using Zemax. In reality, the resolving power of the instrument is determined using collapsed spectra, so the prediction here is on the low side. However, the Zemax analysis assumes perfect optics and alignment. For this reason the predictions have been based on un-collapsed spectra, which is the conservative approach. The expectation is that the actual resolving power will be somewhat better than shown here.

of $1''$ centered in a $2''$ aperture, resulting in a 17.5% throughput loss assumed flat across the bandpass. The telescope efficiency is based on measured data for CO₂-cleaned bare aluminum mirrors (Wilson, RTO II), along with simulated anti-reflection coating curves to match the specifications of the two wide field corrector lenses (a small overall effect). Measured curves were used for




FIG. 27.— The predicted RMS linewidth in pixels as a function of position on the detector for the blue CCDs (left) and the red CCDs (right). The wavelength decreases from top to bottom.

the dichroic, grism, camera coatings, and CCD QE. The manufacturer’s curve for Denton FSS99 silver was used for the collimator coating. A new silvered coating was applied to the collimators in 2012 by Infinite Optics, leading to a system throughput somewhat better than what is described in this paper, particularly at short wavelengths. Internal transmission curves for the camera glasses were obtained from the manufacturer’s data sheets. The fiber efficiency (85%) is based on lab measurements, representing an average value for the fibers measured and assumed flat across the bandpass. An additional 3.6% loss is included for focal ratio degradation beyond the $f/4$ beam the collimator was designed to accept. Figure 28 shows the individual component efficiencies used for this model along with the total expected system throughput. Peak efficiencies in the blue and red channels are similar at about 31%. This represents a major gain relative to the SDSS spectrographs, where a similar component model resulted in expected peak efficiencies of 17% and 22% for the blue and red channels, respectively.

3.4. Mechanical Upgrades

3.4.1. Overview

Figure 29 shows the mechanical layout for the BOSS spectrograph. The revised design is nearly identical to that for the SDSS design. In fact, most of the hardware in the SDSS design is reused for BOSS. Upgrades include new fiber cartridges and slitheads, a new central optics assembly to house new beamsplitters and VPH grisms,

modified lens mounts and hardware to accommodate the revised optical design, and redesigned dewars for the new CCDs. The fiber system upgrades were discussed in Section 3.2. In this section we describe the details of the spectrograph optomechanical upgrades.

3.4.2. Collimator

The collimator opto-mechanics are the same for BOSS as they were for SDSS. However, a subtle, yet important, modification was made by inserting a small mask in front of the collimator. A 3/16 inch wide vertical strip was placed down the middle of the collimator, aligned parallel to the slithead, to prevent reflections off the tips of the fibers and back into the optical system.

3.4.3. Central Optics Assembly

As described in Section 3.3, the right angle grisms used in SDSS were replaced with VPH grisms for BOSS. The new grisms consist of a plane, parallel, VPH grating sandwiched between two prisms. Given the new grism form factor, the central optics assembly was redesigned. The revised configuration can be seen in Figure 29. The new design utilizes the same basic mounting schemes as its predecessor. Each grism is kinematically located by six reference points and registered against those points using spring plungers. Retention of the dichroic is identical to the SDSS design with simple spring plungers to preload the optic against six reference points. As in the original design, all but one of the eighteen locating references (six for each of the three optics) are machined into the base of the assembly, which ensures accurate placement. A mask was added in this design just in front of each grism to prevent undispersed light from sneaking past the grism; something that was discovered during commissioning. Overall the assembly is slightly larger and more massive than the SDSS design. The mass of the BOSS version is 44 kg as compared to 39 kg for the SDSS configuration. Fortunately the opening in the side of the optical bench was large enough to accept the larger-sized assembly. However, a minor adjustment to the guide rails inside the bench was required to allow for a slight shift in the location of the assembly relative to the opening in the bench.

3.4.4. Cameras

Exclusive of the dewars, which are discussed below, only minor opto-mechanical adjustments to the cameras were required to accommodate the BOSS optical design. The BOSS camera designs (both red and blue cameras) have new field-flattening lenses and revised singlet-to-triplet and triplet-to-doublet spacings. To achieve the increased spacing between the singlet and the triplet the singlet cells were remade, the new cells being thicker than the original. The triplet-to-doublet spacings were adjusted by changing the length of the inner barrels. Figure 30 shows a cross-section of the BOSS blue camera.

While the SDSS cameras performed well for many years, the grease coupling layers in the triplet degraded slowly over time, the result being a random web-like pattern of voids, which degraded optical performance. Considerable effort was invested in resolving this problem. Tests using the grease originally used, Dow Corning Q2-3067, showed that these voids begin to appear after just




FIG. 28.— Expected throughput for the BOSS spectrographs. The plot shows all contributors to the throughput model having a wavelength dependence. Not shown are those contributors with an essentially flat response across the bandpass: average fiber transmission including Fresnel losses at the two faces (0.85), focal ratio degradation overfilling the collimator stop (0.96), and “slit” losses for 1" FWHM seeing conditions modeled with a double Gaussian PSF (0.83). Not included in the model are losses due to centering and guiding errors. Overall system throughput, shown by the solid black curve, is expected to peak at about 31% in both the blue and red channels, a significant gain over the original SDSS spectrographs.

one thermal cycle. Hence re-coupling the lenses with fresh grease as a routine maintenance effort was not a viable solution. After much trial and error using alternative optical couplants such as Sylgard 184, and an optical gel, Lightspan LS6943 from Nusil, both of which failed during thermal testing, it was determined that coupling fluid was the only reliable solution.

The triplet cell was redesigned with the lenses being coupled by Cargille Laser Liquid 1074. The revised design is shown in Figure 31. In the new design, the 0.25 mm gap between the lenses is set by a kapton shim with fingers that extend radially inward near the lens outer diameters. The space between the fingers allows the laser liquid to penetrate into the gap. O-rings seal the liquid volume at the rear face of the last element and at the rear face of the front element. Two reservoirs with pistons were integrated as well to accommodate differences in volumetric expansion between the laser liquid and the space it occupies. A drain and fill port, 180° opposed, penetrate the outer diameter of the cell to facilitate filling and draining the liquid.

3.4.5. Dewars

Larger CCDs, larger field flatteners, and new readout electronics necessitated a redesign of the dewars for BOSS. The new design borrows considerably from the original SDSS design with the front half of the dewar containing the field flatteners, the detector, and the pre-amp board, and the rear half of the dewar containing the LN2 reservoir. However, unlike the SDSS configuration, the red and blue detector packages are physically quite different, both from each other and

from the original CCDs. The e2v CCD has a thick silicon carbide package while the fully depleted device from LBNL is mounted to an invar substrate. This led to a redesign of the detector mounting and adjustment scheme, the red and blue designs being unique. The design of the LN2 reservoir is, in fact, nearly identical to the SDSS design. The dewar volume is increased slightly with the usable volume being 0.4 liters. The heat load is 4.4 W, slightly larger than in the original design. With the increased reservoir volume, the hold time is approximately four hours.

3.5. Flexure

Data taken early in the SDSS survey showed discrete spectral/spatial shifts with rotator angle well beyond what was predicted for the bench alone, with the flexure much higher in the dispersion direction than in the spatial direction. For a 360° rotation of the instrument rotator, with the altitude set to the horizon, the flexure in the spectral direction was measured to be 2.2 pixels peak-to-valley for the red camera and 3.6 pixels peak-to-valley for the blue camera. In the spatial direction the flexure was much lower. For the red camera the flexure was measured to be 0.6 pixels peak-to-valley, and for the blue camera 0.5 pixels peak-to-valley. The data also indicated tilt in the focal plane with rotator angle, something that was not anticipated. The fact that the flexure was worse in the spectral direction is sensible given the mounting of the collimator. Mounted behind the center of gravity, the mirror tilts considerably more in the narrow (spectral) direction than in the tall (spatial) di-




FIG. 29.— Cross sections showing the BOSS spectrograph mechanical layout. The layout is identical to the SDSS design, however a revised central optics assembly to accommodate the VPH grisms is implemented here. In addition, there are subtle modifications to the lens group placements within the cameras, and a revised dewar design to accommodate new CCDs.




FIG. 30.— Cross section of the BOSS blue camera. As with the original SDSS design, the red camera opto-mechanical design is identical except for subtle differences in the location of the lens groups.

rection. In addition, shifts in focus as well as hysteresis were noticed in the data. While the flexure was not severe enough to degrade the data in a significant way, the nature of the shifts and the tilt in the detector was troubling. The project lived with the condition for years, and after careful investigation it was finally determined that the camera focus locking mechanism was inadequate, allowing the dewar axis to tilt relative to the optical axis. A redesign of the camera focus locking mechanism was implemented during the initial phase of the BOSS survey, the result being a more stable flexure profile. These results are discussed in some detail in Section 4.

3.6. Detectors, Electronics and Data Acquisition Upgrades

3.6.1. Detectors

The upgrade to 1000, 2" fibers from 640, 3" fibers necessitated the replacement of the CCDs in order to match the fiber diameter to the pixel size. The existing 2k × 2k SITE CCDs with 24 μm pixels were replaced with 4k × 4k CCDs with 15 μm pixels. With some improvements in the optics as described previously, this choice of pixel size provides optimal 3-pixel sampling of the PSF and




FIG. 31.— (a) Exploded view of the BOSS camera triplet, and (b) a cross-section of the BOSS triplet. In this revised design, the lenses are coupled by Cargille Laser Liquid 1074. Lenses are spaced using 0.25 mm Kapton shims with fingers that extend radially inward near the lens outer diameters. Gaps between the fingers allow fluid to flow into the gap between the lenses. O-rings seal the liquid volume at the rear face of the last element and at the rear face of the front element.

the larger format allows for larger gaps between v-groove blocks to sample the wings of the PSF as well as additional pixels in the dispersion direction to accommodate the larger range in wavelength. Equally important, as evidenced in a comparison of Figure 12 to Figure 28, the upgrade to the new devices provides substantially improved quantum efficiency in the regions $3600 \text{ \AA} < \lambda < 5000 \text{ \AA}$ and $7000 \text{ \AA} < \lambda < 10,000 \text{ \AA}$.

3.6.2. Readout Electronics

The readout electronics are almost unchanged from the earlier system, with some additions and modifications required for driving the p-channel LBNL CCDs and handling the larger format of both new detectors. A new dewar board incorporating the preamp (now four channels, one for each quadrant of the new detectors) and clock drivers was designed and built, one version each for the n-channel blue CCDs and the p-channel red CCDs. The OPA627s were replaced with Analog Devices AD8610s, which are slightly noisier but faster and have lower power dissipation; the gain of the new devices is so much higher than that of the earlier CCDs that the slightly higher noise is of no consequence, and the signal chain adds less than a third of an electron in quadrature to the system noise. New and more modern analog switches were incorporated as well, partly to better incorporate the new faster readout (88 kilopixels per second), which is twice as fast as the one used previously. Some of the SDSS camera devices had only one good amplifier and had to be read at this rate, so the signal chain was designed to handle it from the beginning, but better and faster switches make it work better. With the 88 kHz rate, the larger format devices with four channels are read in the same time as the earlier SDSS devices with two.

The dewars are equipped with two hermetic micro-D connectors, one with 51 pins for DC and video, and an-

other with 37 pins for digital clock signals. Cables lead from connectors to saddlebags similar to the originals, but with more boards. The original bus receivers, power-distribution, and signal-chain/bias boards are used with a few gain-setting component changes, but now two sig/bias boards are used, since they were built to handle only two outputs. For the “upside down” LBNL p-channel device, a new board which inverts the bias voltages was designed and built. This board also handles the control of the large substrate bias voltage which places a large DC potential across the device to minimize charge diffusion.

The data and control between the saddlebags and the controller are handled on two cables, a standard 68-pin SCSI cable carrying the RS485 clock signals and a few slow CMOS signals, and a 50-pin MDM cable carrying the digitized video, ADC clocks, and the new voltages.

The controllers are substantially modified. A new power supply to generate the large negative voltages (-30 V) for VDD and VRD and the very high positive substrate bias ($+100 \text{ V}$) for the LBNL devices was added. The single-micro topology was changed to the camera’s two-micro architecture to accommodate the more complex operating environment using two rather different kinds of devices. The clocks are generated synchronously for both red and blue CCDs, but the two sets of clocks themselves are completely different for the two CCD types. The architectures of the detectors are different in detail but also grossly—the E2V blue devices are 4-phase in the parallels, but the LBNL devices are three-phase, and the way charge is passed to the serial register is rather different.

The FOXIs have been replaced by a fast ethernet-interfaced micro, the Netburner, described in the next subsection.

3.6.3. Data Acquisition

The data acquisition (DAQ) system is responsible for data transfer from the spectrograph to the instrument control computer (ICC), located approximately 300 feet from the telescope. The DAQ receives pixel data from the spectrograph electronics, performs some processing including simple data integrity tests, formats the data according to the FITS standard, and saves the images on a disk array. These operations have to be performed in real time to maximize throughput and overall survey efficiency. The architecture of the BOSS DAQ is similar to the proven design of the original SDSS DAQ, described in Section 2.5. The main change was to use standard commodity computing components to replace the original, aging hardware modules, many of them custom-built and no longer supported. An additional benefit was the replacement of licensed software packages with free, open-source products.

A block diagram of the BOSS DAQ system is shown in Figure 32. Separate DAQ modules connect to the digitizer boards of the two cameras in each of the spectrographs. Copper Ethernet cables connect the DAQ modules to a network switch mounted on the telescope, where the data is converted to optical signals and travels over optical fibers to the computer housed in a separate building.




FIG. 32.— Block diagram of the upgraded BOSS DAQ system.

DAQ Hardware— To efficiently collect, package, monitor, and transmit data from the BOSS spectrographs, a custom printed circuit card was designed. The card, called the DAQ board, is shown in Figure 33, and a block diagram of its contents are shown in Figure 34. One DAQ board is required for each BOSS spectrograph. During normal operation, the DAQ board receives 8 synchronous streams of pixel data, coming from the two cameras on each spectrograph, each camera having one CCD with four outputs. The incoming data, in RS-485 format, are stored in a frame buffer to decouple input and output operations, reformatted and transferred via a 100 Mbit Ethernet link to the ICC over optical fiber to the computer room. The input stage of the DAQ board operates at the pixel clock rate set by the front-end electronics. The buffer memory is of sufficient size to store an entire exposure, but under normal circumstances the network interface is fast enough to keep up with the input rate. In order to fit in the limited space available on the spectrograph, the DAQ card was designed to be compact, 2.5×4.5 inches, and to operate from a single 5 V supply. The design of the DAQ card is based on three main components; a Xilinx FPGA, a NetBurner Ethernet Core Module, and eight $8M \times 16$ pseudo-static ram chips (PSRAM). An overview of these components and their functions is given below.

The Xilinx FPGA, a Spartan XC3S1400AN, controls the input stage of the DAQ module and the frame buffer memory. The bus interface to the NetBurner module is also implemented in the FPGA. The pixel data streams are de-serialized and recorded in the frame buffer which is organized as a First-In/First-Out (FiFo) memory. The FPGA is configured to write to all eight memory modules simultaneously, thus maximizing the data recording speed. At the FPGA level data integrity checks are performed by keeping pixel and line counts. The acquisition software running on the NetBurner module can access this information and, if necessary, raise appropriate flags in the ICC. For testing and debugging purposes the FPGA also contains logic to emulate the spectrograph signals and load test images to the frame buffer, which in turn can be read and sent to the ICC. With this feature, problems upstream from the DAQ card may be




FIG. 33.— The BOSS DAQ custom printed circuit board, top and bottom sides.




FIG. 34.— Block diagram of the BOSS DAQ card.

located and the ICC and analysis code can be debugged easily using well-defined images. In addition, the FPGA firmware includes a watchdog timer that will raise a flag if the camera data is not loaded into the DAQ card within a reasonable time limit.

The NetBurner Ethernet Core Module, model 5234, coordinates the communication between the DAQ card and the ICC. The software for the NetBurner is written in C/C++ and will be discussed in detail in the following section. The NetBurner contains a Freescale ColdFire 5234 microprocessor, 2 MB Flash memory, 8 MBytes of SDRAM, and 10/100 Ethernet interface. The ColdFire microprocessor has a full 32-bit architecture providing more than 140 million instructions per second of processing power. The NetBurner module is interfaced to the FPGA via a 16-bit external data bus which provides a simple and fast communication link.

The $8M \times 16$ pseudo-static DRAM ram chips used to buffer and store the spectrograph images are Micron MT45W8MW16BGX. Pseudo-static PSRAM is a memory technology that features an SRAM-like architecture using a DRAM cell with internal refresh operations that do not interfere with user read or write cycles. This approach significantly reduces the complexity of the DAQ

board while still allowing for sufficient memory to buffer an entire exposure. Should a problem arise with the communication link to the ICC, this feature allows the data to be resent repeatedly until it is received correctly. The PSRAM type chosen is limited to a read/write timing of 70 ns, but consumes little power; less than 50mA, during a read or write cycle.

DAQ Software— The DAQ software consists of two major pieces. The acquisition part, written in C, runs on the NetBurner modules on each of the DAQ boards. The second software component, a Python module, provides image formatting as well as an interface to the BOSS online system. This software is executed on the instrument control computer in the computer room. The acquisition program takes advantage of the soft real time functionality provided by the NetBurner software tools. One task implements a command interface and support for the emulation mode used for hardware and software debugging. A second task is responsible for reading out the data from the frame buffer on the DAQ module and for data transmission to the ICC using the TCP/IP protocol. Additional functionality provided by the application includes a separate reset connection to restart the software in case of problems, as well as telnet support to allow remote connections to the NetBurner low-level debugger. Furthermore, the NetBurner allows for software updates over TCP/IP, permitting remote modification and testing in the event that the acquisition software developer is not on site.

Both BOSS spectrographs operate independently at the DAQ level and are seen by the online system as independent devices. A Python module was developed to allow the online system to interact with the DAQ controllers. Responsibilities for the data transfer, connection management and data integrity checks, such as the correct number of pixels, lie with this module. The Python module also rearranges the pixel streams from the four amplifiers for each of the two CCDs, generating FITS formatted files with contiguous data sections and with the overscan regions moved to the edges of the array. Header keywords are generated and added and for each camera a separate FITS file is written to disk.

DAQ Performance— In almost two years of survey operation the BOSS DAQ system has performed efficiently and reliably. The readout time has been measured to be 55.6 seconds. Only 1.6 seconds of this period is added by the DAQ system, while the bulk of the readout time is used for CCD digitization and is set by the clocking speed, which is optimized to minimize readout noise.

3.7. *Guider Upgrade*

3.7.1. *Overview*

As mentioned in Section 2.7, the original Roper Scientific camera used for the guider in SDSS-I and -II was unsatisfactory on several counts, and it was decided at the outset of SDSS-III to replace it with a cooled camera with a better detector. In addition, although the pointing of the telescope is usually accurate enough that guide stars are found in two large (11") guide fibers, the efficiency of the survey would be enhanced if we had larger acquisition fibers, and we implemented two of these per cartridge in SDSS-III.

With the original SDSS system, focus was achieved by trial and error, adjusting focus of the telescope and tracking changes in the FWHM of the stars in the guide fibers. The focus system for the SDSS *imaging* camera worked quite well. In this system, stars were simultaneously imaged onto one focal plane that was set some distance behind the scan data and another set of stars were imaged onto a focal plane that was set some distance in front of the scan data. From the images mapped in these three focal planes, it was possible to derive a focus error signal, and it was decided to implement a similar design in the new spectroscopic guide system for SDSS-III. In a similar fashion to the imaging system, a set of guide fibers was placed in the same focal plane as the science fibers for the spectrograph, a second set was slightly displaced outside of focus, and a third set inside. In addition to the two large acquisition fibers, the number of guide fibers was increased from 11 to 14 to account for any potential decrease sensitivity in the unfocused fibers to perform the other guider functions.

3.7.2. *Guider Camera*

The new guider camera is an Alta E47 from Apogee Imaging Systems³². The camera employs a 100 Mb/s ethernet interface, uses a thermoelectrically cooled e2v CCD47-10 AIMO 1024 × 1024 back-illuminated CCD with 13.5μm pixels, normally binned 2×2. This results in 27 μm, 0.45" pixels. The images are nearly always well-sampled for the seeing conditions at APO. The camera is much more sensitive and stable than the old one, and allows the use of much fainter guide stars.

The camera was supplied with a liquid cooling head, which we replaced with a similar part made of copper. The replacement is intended to alleviate prior bad experiences in the glycol system with aluminum parts. Otherwise, the camera is used as supplied. The reimaging system from the output fiber block to the CCD and the broad visual (Bg38) filter from the old system was retained for the new guider.

3.7.3. *Guide Fibers*

The new system uses 14 7" (400 micron) Sumitomo coherent fibers of the same kind used in the old system for guiding. Eight of these are in ferrules which place them at the nominal focus of the plate, and three each are 200 microns inside and outside of nominal focus. The focus offset adds 0.7" FWHM in quadrature to the seeing of the images, which is not yet enough to make them unsatisfactory for guiding, but still provides a sensitive measure of focus.

In addition to the 14 guide fibers, there are two much larger (24") acquisition fibers, which are inexpensive coherent plastic fiber bundles, used for guide star acquisition only. Though not used not for guiding, their quality seems to be almost as good as the glass guide fibers.

3.7.4. *Software, Control, and Performance*

The software for the new control system is an adaptation of the old guider code, written mostly in C. There are a few enhancements over the old system, including a fast acquisition mode that uses the acquisition fibers

³² <http://www.ccd.com/>

only, a mode in which successive guider images can be accumulated to improve stability when using bright guide stars and therefore short exposures, and a slow servo loop to actively and automatically control the scale of the telescope (in the SDSS-I/II surveys the scale of the telescope was adjusted manually). The upgraded BOSS guider automatically tracks in RA, Dec, rotation, and scale. All of the fits used to generate guider offsets are simple least-square deviation fits; it is possible that we could do better if a refraction model using the actual distribution of the guide stars on the plate were used, but this has never been implemented. The present system is satisfactory if the guide stars are reasonably uniformly distributed over the plate; we strive to achieve this distribution in the process of designing each plate.

The display code was completely rewritten for the new system but is functionally identical to the old one. A ‘raw’ CCD image is displayed to the observers as is an image in which the fiber images are distributed approximately as they are on the sky. This display makes qualitative evaluation of the guider behavior easy. In addition, the guider errors are displayed for each fiber for each image and are recorded in the guider frame headers.

4. SDSS AND BOSS SPECTROGRAPH PERFORMANCE

SDSS and BOSS observations of each plate follow the sequence: mount cartridge, slew to field, focus spectrograph, take exposure with calibration arc lamps and flat field lamps, and obtain a series of 15 minute science exposures. More details on the procedures of spectroscopic observations can be found in Dawson et al. (2012). Here we describe the methods for determining the instrument performance for SDSS and BOSS. Instrument characterization is presented in three parts: the spectroscopic pipeline, the spectrograph performance, and CCD performance. Data found in DR8 (Aihara et al. 2011) is used in the characterization of the SDSS spectrograph and data in DR9 (Ahn et al. 2012) is used in the characterization of BOSS.

4.1. Spectroscopic Pipeline

The spectroscopic redshifts and classification for SDSS and BOSS are based on the data reduction pipeline **idl-spec2d**³³. A brief explanation of the routines for SDSS is found in Adelman-McCarthy et al. (2008) and Aihara et al. (2011). The pipeline has been modified for BOSS to account for the increase in wavelength coverage, number of fibers, and other changes discussed in Section 3. For a full description of the software reductions, see Schlegel et al. (2012) and Bolton et al. (2012). The intermediate data products produced by the **idl-spec2d** pipeline are used to characterize the instrument performance; we include a brief explanation of the relevant functions here.

In the first step of the data reduction, the read noise and bias level from the raw images is determined from the overscan of the image. Next, raw images are bias-subtracted and counts are converted into electrons using CCD amplifier gains specific to each quadrant of each detector. Cosmic rays are then identified using a filter to discriminate between the sharp edges of cosmic ray trails and the smoother profile of optical sources. The

noise of each pixel is determined by taking the Poisson noise from the electron counts in quadrature with the measured readnoise. Finally, a flat-field image is used to correct the pixel-to-pixel response in the CCD.

Each fiber projects a spectrum onto the CCD with an RMS in the spatial direction of roughly one pixel and a spectral dispersion of roughly one pixel RMS in the direction of the parallel clocking. The data from the flat field exposures are used to model the profile of each fiber on the CCD and also to normalize fiber-to-fiber throughput variations. The data from the arc exposures are used to determine the wavelength solution for each fiber.

Figure 35 shows an example of the arc lines obtained in a BOSS exposure. The top left panel is taken near the edge of the blue CCD and shows a slight tilt in the wavelength direction due to optical distortion. This image also demonstrates the bundling of 20 fibers separated by a gap on either side. The top right panel demonstrates a dropped fiber near the right edge of the first bundle. The bottom left image from a red camera is taken near the center of the CCD and shows the amplifier gain between two quadrants at a different background count. The lower right panel shows several cosmic rays interacting with the CCD on a red camera. The bottom right image also shows an offset in the wavelength direction between the central bundle and the surrounding bundles due to imperfections in the slithead assembly. The top two images are more representative of the SDSS CCDs because the red BOSS CCDs are much thicker and lead to more pronounced cosmic ray trails. We use the wavelength solution and line spread function derived from the arc exposures to determine wavelength coverage and resolution.




FIG. 35.— Raw arc image taken with the BOSS spectrograph. Top left panel: 3650 – 3850 Å on blue CCD; Top right panel: 5050 – 5250 Å on blue CCD; Bottom left panel: 6000 – 6250 Å on red CCD; Bottom right panel: 8850 – 9100 Å on red CCD. Each panel covers roughly 50 spectra.

³³ Current and development versions of the software are found at <http://www.sdss3.org/dr9/software/products.php>

Each two dimensional science image is collapsed into a

series of one-dimensional spectra using the fitted profile and wavelength solution from the calibration exposures for each fiber. At this point, the “SKY” fibers (fibers assigned to areas with no detected objects) are used to model the background as a function of position on the detector. The model sky background is subtracted from the one-dimensional spectra. These one-dimensional spectra contain all of the information required to characterize the instrument performance. For science applications, the spectrum from each fiber is resampled to a pixel size of 69 km/s and combined with the spectra from the same fiber for each exposure in the sequence to generate a high signal-to-noise spectrum. Redshifts and object classification are determined from these co-added spectra as will be described in Bolton et al. (2012).

SDSS data were processed using `idlspec2d v5.3.12`, which is the same version of the software used in DR8. BOSS data were processed using `v5.4.31`, a somewhat earlier version of the `v5.4.45` software used for DR9 (Ahn et al. 2012). There are minor changes between these two versions, but the differences do not change the conclusions of the work presented here.

4.2. Spectrograph Performance

4.2.1. Wavelength Coverage

The SDSS and BOSS CCDs were described in Section 2.5 and Section 3.6, respectively. Because the optics are relatively similar, the fundamental difference between the two surveys is the size of the detector and the resulting wavelength coverage. The new BOSS CCDs are 6 cm on a side, compared to the 4.8 cm SDSS detectors. The difference results in a substantial increase in the BOSS wavelength coverage.

Besides the obvious motivation to use recent advances in CCD fabrication to increase the detector size for BOSS, the scientific rationale behind the wavelength coverage for the two surveys was described in Section 2.1 and Section 3.1. The exact measured wavelength coverage is determined from the range of arc lines detectable in the arc exposures for the two surveys. The highest energy arc line detected in SDSS calibration exposures is Hg I at 3901.87 Å, ensuring that CaII H and K lines are detected at zero redshift. The lowest energy arc line detected in SDSS calibration exposures is Ne I at 9148.67 Å, allowing H- α to be detected out to a redshift $z = 0.4$. The BOSS blue limit includes the wavelength-calibration Cd I arc line at 3610.51 Å, providing extended coverage of the Lyman- α forest lines relative to SDSS. The BOSS red limit includes the detection of a Hg line at 10140 Å, covering the Mg b 5175 Å and the Na D 5893 Å absorption features redshifted to $z < 0.72$. In between these detected arc lines for SDSS (3902 Å – 9149 Å) and BOSS (3611 Å – 10140 Å), we find a reliable wavelength solution. The wavelength coverage for SDSS and BOSS leads to detection of at least 31 common astrophysical emission lines, as detailed in Table 5 of Bolton et al. (2012).

4.2.2. Spectral Resolution

The upgrades to BOSS described in Section 3 do not significantly impact the spectral resolution of BOSS compared to SDSS. The resolution requirement for SDSS was broadly defined to balance the tradeoff between resolving absorption lines in relatively massive galaxies and

maximizing wavelength coverage. The resolution requirement for BOSS is driven by the ability to fit the Balmer series between $3800 < \lambda < 4900$ Å. From these fits we model surface gravity and stellar parameters and derive a synthetic spectrum for each standard star used in spectrophotometric calibration for the exposure. Experience with SDSS indicates that a resolving power lower than $R = 1400$ in this wavelength range was inadequate for precise and reliable stellar classification. The BOSS galaxy spectra require resolution sufficient for achieving RMS redshift error of 300 km/s, achieved with a resolving power of $R > 1000$ for the remainder of the wavelength range.

The spectral resolution is measured from calibration arc images taken before each set of science exposures. The one-dimensional arc image is first masked to include only pixels within 12 pixels of the center of each arc line on each fiber. A Gaussian of width σ_λ is fit to each spectral profile using the 25 unmasked pixels. A fourth order Legendre polynomial model is fit to the derived σ_λ as a function of wavelength to model the dispersion over the full wavelength range. The resolving power for SDSS and BOSS is then formally defined as $R = \frac{\lambda}{2.35 \times \sigma_\lambda}$. The *resolution* is defined as the FWHM of the Gaussian, $2.35 \times \sigma_\lambda$.

We measured the resolving power as a function of wavelength from a sample of 100 SDSS plates and 100 BOSS plates. For each plate and each camera, we computed the mean R as a function of wavelength. Figure 36 shows a comparison of the SDSS resolving power to that for BOSS.

The presence of a dichroic for SDSS and BOSS is revealed near 6000 Å where the resolving power no longer increases monotonically with wavelength. The extended coverage at long wavelengths from the new BOSS CCDs is evident redward of 9200 Å while the extended coverage at short wavelengths is more subtle, but also shown in the figure.

The BOSS requirements and measured resolving power are presented in Figure 37. For BOSS, the position of the fiber on the CCD impacts the resolution due to optical distortion over the larger detectors. The resolution of the middle fiber is representative of $\sim 80\%$ of the BOSS fibers, 100% of the SDSS fibers, and clearly exceeds specifications at all wavelengths. As the fibers get closer to the edge of the spectrograph, the resolving power decreases as predicted in Section 3.3. A small fraction ($\lesssim 10\%$) of the BOSS fibers therefore do not meet specifications for spectral resolving power in the wavelength range $3800 < \lambda < 4900$ Å.

In addition to the resolution as a function of wavelength, we have also determined the RMS width σ_p of the line-spread function in native pixels to demonstrate the spatially varying resolution over the focal plane of the BOSS CCDs. Figure 38 shows the RMS pixel resolution of the illuminated region for each of the four BOSS CCDs. For the blue cameras, the spectra near the edges of the detector have a higher σ_p , reflecting the lower optical quality and deviations from flatness of the focal plane at large field angles. The optics project a saddle-shaped focal plane which contributes to this poorer resolution as can be seen in the upper panels. The two red cameras display increased RMS width between $9500 < \lambda < 10200$




FIG. 36.— The resolving power for BOSS (gray) and SDSS (red). The shaded regions correspond to the regions that contain 68% of the plates in the measurement. The top panel shows the results for Spectrograph 1 and the bottom panel shows the results for Spectrograph 2.

Å, mostly as a consequence of the increased path length for the generation of electron-hole pairs in the silicon in the thick CCD at longer wavelengths. Also shown is the region corresponding to the SDSS CCDs which are not subject to these effects on either the blue or red cameras.

4.2.3. Throughput

As discussed in Section 3, the primary motivation for rebuilding the spectrographs for BOSS was to increase the instrument throughput. We define the throughput as the ratio of the measured flux for a point source relative to the incoming flux outside the atmosphere. A higher throughput is important for the fainter objects targeted by BOSS, particularly the highest redshift galaxies and the quasars near the magnitude limit of the survey. The throughput described here is the throughput of the entire instrument; the predicted throughput of the individual components was discussed in Section 2.3 and Section 3.3.

The throughput is measured from the spectrophotometric standard stars on the plate during science exposures. For each standard star, raw photon counts are measured using a 3-pixel radius boxcar extraction on the calibrated, two-dimensional science frame before flux calibration. The aperture is centered on the spectral trace using the fiber flat-fields from the calibration sequence. The photon counts from the source are estimated at each pixel using the observed magnitudes from the SDSS imaging survey and the synthetic spectral tem-




FIG. 37.— The requirements and the measured resolution for the two sets of BOSS spectrographs. The dashed black curve is the requirement for the resolution. The green curve is the 68% confidence limit about the mean of the resolution for the central fiber and is representative of $\sim 80\%$ of the fibers. The gray curve is the 68% confidence limit about the mean of the resolution for a representative fiber near the edge of the spectrograph.

plate derived from spectrum. The synthetic spectrum is converted to photons per pixel using the wavelength solution, the effective collecting area of the obstructed primary mirror, and the time of each exposure. The throughput at each pixel is calculated by taking the ratio of the measured raw counts to the number of photons per pixel as predicted by the model.

The throughput for SDSS was measured by averaging the throughput from 84 different standard stars observed under photometric conditions at an airmass of ~ 1.0 and seeing $\leq 1.15''$. Because BOSS has smaller fibers, the measured throughput on the BOSS spectrographs is more susceptible to guiding errors. We attempt to mitigate this effect by averaging the throughput over the four stars that produce the highest throughput on each camera. In both cases, we correct small differences in seeing by normalizing the observed flux to a Gaussian with FWHM of $1''$ integrated over an aperture that corresponds to the area of the $2''$ or $3''$ fiber. The throughput curves have been flat-field corrected to account for fiber-to-fiber variations, thereby normalizing each fiber to the median for each plate.

As shown in Figure 39, the throughput after the BOSS upgrades has been significantly improved. The ratio of BOSS to SDSS throughput is also shown in Figure 40. Not demonstrated here is a fiber dependent variation




FIG. 38.— The RMS width in pixels in the wavelength direction for each of the four BOSS CCDs. The top panels represent the two blue cameras while the bottom panels represent the two red cameras. In all cases, the wavelength increases from the bottom of the image to the top of the image. The SDSS images have similar characteristics and are not shown here.




FIG. 39.— Throughput curves for SDSS (black) and BOSS (red). The results for Spectrograph 1 are displayed in the top panel. The results for Spectrograph 2 in the bottom panel. Note the telluric absorption features.

in the throughput. As in the resolution measurements, fibers near the edge of the BOSS detectors have a slightly lower throughput than fibers near the middle of each spectrograph. For BOSS, there is an additional $\sim 10\%$ RMS variation on the average throughput curve (particularly at short wavelengths) between standard stars. This variation in BOSS is likely caused by scale changes in the telescope over the course of an observation and possible guiding errors (BOSS is more sensitive to guiding because of the smaller fibers). The effect is probably present in the SDSS data as well, but is amplified in BOSS due to the decrease in fiber diameter.




FIG. 40.— Ratio of throughput BOSS/SDSS after applying a median smoothing kernel of width 100 pixels. The figure displays results for Spectrograph 1 (black) and Spectrograph 2 (red).

4.3. Flexure

With the implementation of a redesigned focus locking mechanism for the BOSS cameras, the gravity induced image motion behaves in a predictable and stable way; see Figure 41. The figure shows flexure as a function of rotator angle for BOSS spectrograph #2. To generate these plots, data was collected with the telescope zenith angle at 60° , then scaled to a zenith angle of 90° ; the worst case pointing scenario. The sign convention is such that positive rotator angles represent a counterclockwise rotation of the rotator as viewed from behind the rotator. Positive values in the spectral direction represent motion of the spectra toward longer wavelengths, and in the spatial direction positive values represent motion of the spectra toward higher column number on the CCD. Flexure is largest in the spectral direction due, in part, to the form-factor of the collimator. The support points behind the tall mirror are closer together in the spectral direction; hence motion, in angular terms, due to compliance in the actuators is larger. For a one hour exposure, at a zenith angle of 60° , the worst-case spectral shift is ~ 6 microns, about 0.4 pixels on the BOSS CCDs.

4.4. CCD and Electronics Performance

In addition to the spectrograph performance, we also measured the CCD and electronics performance for SDSS and BOSS. First, the CCD images were checked for cosmetic defects such as saturated hot pixels, pixels with a large deviation from neighbors, blocked columns at low flux levels, and other bad pixels. These bad pixels are masked during the data processing. The hot pixels are recorded by identifying any pixel greater than 15,000 ADU in the bias exposures. Hot columns are identified in bias exposures as regions of connected pixels that lie above the background, in the direction of parallel clock transfer, by an amount equal to five times the readnoise. Pixel-to-pixel variation was mapped using a lossy-fiber that was placed in the position of the usual fibers along the slithead of an engineering cartridge. While the usual flat fields are illuminated through discrete fibers, the lossy-fiber spectra follow a similar spectral dependence but have nearly uniform illumination in the spatial direction. Blocked columns and defective pixels are identified in these images as clustering of pixels that lie 50% below the mean background. The flat fields from these tests are also applied at the same time as the bad pixel masks in the data processing. Other bad pixels are found from a combination of flats, biases, and dark frames.

In addition to cosmetic defects, the read noise, dark current and gain of the CCDs were measured. The read noise was measured in a series of zero second exposures; bias frames that were obtained throughout the surveys. To maintain Poisson-limited statistics in the region of Lyman- α forest, the detectors in the BOSS blue cameras require a read noise less than $3.0 e^-/\text{pixel}$ RMS. The requirement for the red cameras is less stringent because the spectra are sky-dominated for most of the wavelength coverage on the red cameras. The read noise for the red cameras is required to be less than $5.0 e^-/\text{pixel}$ RMS. The read noise requirement is met for each CCD and can be found in Table 3.

The measurements of dark current on the BOSS detectors are shown in Table 3. To prevent significant




FIG. 41.— Gravity induced image motion in BOSS spectrograph as a function of rotator angle for a zenith angle of 90° . The top and bottom plots show flexure in the red and blue channels, respectively. Flexure in the spectral direction is greater than in the spatial direction due to compliance in the collimator mount; a natural consequence of a tall mirror.

contribution to the noise in the measurement, we require a dark current on the BOSS detectors less than $1.0 e^-$ for the blue CCDs and less than $2.0 e^-$ for the red CCDs during a 15 minute exposure. We include any light source detected by the CCDs that does not originate from the fibers, including light leaks in our measurement of dark current. A series of 15 minute exposures is averaged pixel-by-pixel to determine the dark current in each pixel. The dark current in each quadrant of a CCD is averaged, and we report the minimum and maximum values of dark current in the four quadrants in Table 3. The read noise and dark current requirements are particularly important at short wavelengths where there are low sky and photon counts. The gain is used in determining errors assigned to each pixel.

For BOSS, we use measurements in the detector lab at LBNL to determine full-well depth for the red CCDs. We refer to documentation from e2v for these measurements on the blue CCDs. The full-well depth exceeds the maximum number of electrons expected ($65,000 e^-/\text{pixel}$) in

TABLE 3
CCD PERFORMANCE FOR SDSS AND BOSS

	read noise (e^-)	dark current ($e^-/\text{pix}/15 \text{ min}$)	gain (e^-/ADU)	number of bad columns	fraction of bad pixels
SDSS					
b1	2.8-3.8	1.05-1.10	1.05-1.10	4	
b2	3.2-3.9	1.23-1.26	1.23-1.26	4	
r1	3.5-4.1	1.00-1.09	1.00-1.09	14	
r2	3.6-4.4	1.04-1.05	1.05	10	
Start of BOSS					
b1	1.79-1.98	0.64-0.68	1.01-1.05	0	2.0e-05
b2	1.74-2.04	0.61-0.65	0.99-1.04	2	2.1e-04
r1	2.36-2.72	0.88-1.10	1.54-1.97	5	1.4e-04
r2	2.42-2.99	0.95-1.97	1.54-1.96	3	1.8e-04
MJD 55300 (r2 replaced)					
b1	1.83-2.01	0.51-0.53	1.01-1.05	0	2.1e-05
b2	1.87-2.03	0.53-0.56	0.99-1.04	2	2.1e-04
r1	2.43-2.73	0.63-0.80	1.54-1.97	5	3.3e-04
r2	2.73-2.89	1.19-2.27	1.59-1.66	11	4.4e-04
MJD 55800 (r1 replaced)					
b1	1.77-2.02	0.46-0.49	1.01-1.05	0	2.1e-05
b2	1.86-2.01	0.56-0.59	0.99-1.04	2	2.1e-04
r1	2.45-2.82	0.57-0.82	1.47-1.93	1	3.3e-04
r2	2.85-2.88	1.30-1.56	1.59-1.66	11	2.5e-04

a science exposure for all cameras. The charge transfer efficiency (CTE) records the fraction of electrons that are successfully transferred during each clock cycle during readout. CTE for the red CCDs was measured using an ^{55}Fe source (as described in Bebek et al. 2002; Dawson et al. 2008) in the LBNL detector lab for the red CCDs and recorded from the e2v documentation for the blue CCDs. All detectors exceed the CTE requirement of 0.99999 in both serial transfer and parallel transfer. Finally, to meet our goal for overall observing efficiency, a maximum of 70 seconds can be used for the CCD pre-exposure flush time and end-of-exposure readout time. The readout time of the BOSS CCDs has been measured to be 55.6 seconds as described above. This requirement means that not more than 10% of the observing time is used for the CCD flush and readout for two calibration exposures and four 900 second science exposures. Because the detectors have been replaced by newer technologies, we do not report values for full-well depth, CTE, or readout time for the SDSS detectors.

5. CONCLUSION

The SDSS spectrographs were designed to be high-throughput, robust instruments, to conduct an extragalactic survey producing a million redshifts over a five year lifetime. These goals have clearly been met, with 1.6 million spectra of stars, galaxies, and quasars over 2880 plates, primarily in the northern Galactic cap at high Galactic latitudes. Spectra over 9274 deg^2 were collected in nine years of operation between the start of SDSS-I in 1999 and the completion of the first of two extended phases, SDSS-II, including SEGUE which produced roughly 250,000 stellar spectra between 2004-2008 (Yanny et al. 2009). These data were released to the astronomy community nearly every year; the full data set from SDSS-I and II is included in DR8 (Aihara et al. 2011). Minor technical issues aside, the spectrographs proved to be quite reliable and exceptionally productive instruments. With improvements in CCD technology,

the use of VPH gratings, and enhancements to the optics (i.e. fluid coupled camera triplets, new dichroics, and higher reflectivity collimator coatings) the instrumental throughput was enhanced significantly, allowing the use of smaller diameter fibers and making the BOSS survey possible. Along with improvements to the guide camera, modifications to improve flexure, and a host of minor upgrades, the BOSS spectrographs have proven to be as productive as the original SDSS design and have many years of life remaining.

The spectroscopic data from SDSS have enabled studies across a broad range of astronomical disciplines including the evolution and clustering of galaxies (e.g., Kauffmann et al. 2004; Tegmark et al. 2004), gravitational lensing (e.g., Bolton et al. 2006), the properties of quasars (e.g., Vanden Berk et al. 2001), and stellar astrophysics (e.g., West et al. 2008). One of the prominent scientific contributions from SDSS and SDSS-II data is the discovery of acoustic oscillation signatures in the clustering of galaxies (Eisenstein et al. 2005), opening the door to a new method of cosmological measurement. SDSS imaging and spectroscopic data have been included in more than 3500 refereed papers.

As described in detail in Dawson et al. (2012), the BOSS survey will cover $10,000 \text{ deg}^2$ in five years, obtaining spectra for 1.35 million luminous galaxies with redshifts $0.15 < z < 0.7$, and approximately 160,000 quasars with redshifts between $2.15 < z < 3.5$. The galaxies provide a highly biased tracer of matter, and at these densities over this area provide percent-level distance scale measurements through studies of BAO. The quasars provide a backlight to illuminate neutral hydrogen through Lyman- α absorption along the line of sight, thereby mapping large-scale structure in the foreground of each quasar. These distance measurements will represent a significant improvement on the accuracy of existing BAO measurements, and will tighten the constraints on the dark energy equation of state, $w = p/\rho$, where p is pressure and ρ is density, and the evolution of w with

time. Detection of the BAO feature in quasar spectra would allow BOSS to make the first characterization of dark energy at early times, when dark energy should be sub-dominant according to the prevailing Λ CDM theory.

BOSS completed imaging of the SGC spectroscopic footprint in 2008 and 2009; spectroscopic observations began December 5, 2009 after a three month commissioning phase. The algorithm for selecting galaxy targets is outlined briefly in White et al. (2011) and Eisenstein et al. (2011), and described in detail in Padmanabhan et al. (2012). The quasar target selection is described in Ross et al. (2012) using photometry from the SDSS imaging survey supplemented with data from programs at ultraviolet, infrared, and radio wavelengths to enhance the detection efficiency. These first two years of BOSS spectra include 324,198 unique galaxy targets ($z > 0.43$) with a 98.7% rate of successful classification, and 103,729 unique lower redshift galaxy targets ($0.15 < z < 0.43$) with a 99.9% successful rate of classification. There were more than 85,000 quasars observed, 61,933 of which were at $z > 2.15$. The spectra of these objects are classified as described in Bolton et al. (2012) and each quasar is visually inspected and recorded as explained in Pâris et al. (2012). Data from the first two years of the survey were included in DR9 (Ahn et al. 2012) and include the spectra from 831 plates.

The higher redshift sample of galaxies from DR9 has already been used to fit the position of the acoustic peak at an effective redshift $z = 0.57$ (Anderson et al. 2012). With 1.7% accuracy, this is the most precise distance constraint ever obtained from a galaxy survey at any redshift. The potential for distance measurements at $z > 2$ using Lyman- α forest absorption from quasars is promising; the first detection of flux correlations across widely separated sightlines was obtained in the first year BOSS data to comoving separations of $60 h^{-1}$ Mpc (Slosar et al. 2011). BOSS is on pace to complete the 10,000 deg² spectroscopic footprint by the middle of 2014, providing new cosmology constraints and a wealth of data for studies of galaxy evolution and quasar physics.

The optical spectrographs at APO will have unique capabilities for widefield, multiplexed spectroscopy well beyond the 2014 completion of the BOSS survey. The SDSS collaboration has proposed an “After SDSS III” (AS3) program for 2014-2020 that will capitalize on this

resource. Four separate programs in AS3 will use the BOSS spectrograph, likely with only minor modifications. The extended Baryon Oscillation Spectroscopic Survey (eBOSS) will probe dark energy and fundamental physics by making distance measurements with BAO in the redshift range $0.6 < z < 2$. The Time-Domain Spectroscopic Survey (TDSS) will obtain spectra of 100,000 Galactic and extragalactic variable sources with simultaneous, multi-epoch optical imaging data. In a program to follow up x-ray sources in new 0.28 keV data obtained from the extended ROentgen Survey with an Imaging Telescope Array (eROSITA: Predehl et al. 2010), the SPectroscopic IDentification of eROSITA Sources (SPIDERS) survey will follow up 50,000 – 100,000 objects. Finally, Mapping Nearby Galaxies at Apache Point Observatory (MaNGA) will perform spatially resolved spectroscopy on approximately 10,000 nearby galaxies using 15 integral field Units integrated into new BOSS-like cartridges. As with the original SDSS spectroscopic survey, these four data surveys will provide a premier data sample for astrophysical studies from Galactic to cosmological scales.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

REFERENCES

- Abazajian, K. N., et al. 2009, *ApJS*, 182, 543
 Adelman-McCarthy, J. K., et al. 2008, *ApJS*, 175, 297
 Ahn, C., et al. 2012, *ArXiv e-prints*
 Aihara, H., et al. 2011, *ApJS*, 193, 29
 Anderson, L., et al. 2012, *ArXiv e-prints*
 Bebek, C., et al. 2002, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4669, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. M. M. Blouke, J. Canosa, & N. Sampat, 161–171
 Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, *ApJ*, 638, 703
 Bolton, A. S., et al. 2012, *ArXiv e-prints*
 Burgh, E. B., Bershad, M. A., Westfall, K. B., & Nordsieck, K. H. 2007, *PASP*, 119, 1069
 Cohen, J. G., Oke, J. B., Carr, M., Harris, F. H., & Hamilton, D. 1988, in *Astronomical Society of the Pacific Conference Series*, Vol. 3, *Fiber Optics in Astronomy*, ed. S. C. Barden, 190–194
 Dawson, K., et al. 2008, *IEEE Transactions on Nuclear Science*, 55, 1725
 Dawson, K. S., et al. 2012, *ArXiv e-prints*
 Eisenstein, D. J., et al. 2001, *AJ*, 122, 2267
 Eisenstein, D. J., et al. 2011, *AJ*, 142, 72
 Eisenstein, D. J., et al. 2005, *ApJ*, 633, 560
 Falco, E. E., et al. 1999, *PASP*, 111, 438
 Fukugita, M., Ichikawa, T., Gunn, J. E., Doi, M., Shimasaku, K., & Schneider, D. P. 1996, *AJ*, 111, 1748
 Gunn, J. E., et al. 1998, *AJ*, 116, 3040
 Gunn, J. E., et al. 2006, *AJ*, 131, 2332
 Holland, S. E., et al. 2006, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 6276, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
 Huchra, J., Davis, M., Latham, D., & Tonry, J. 1983, *ApJS*, 52, 89
 Ivezić, Ž., et al. 2004, *Astronomische Nachrichten*, 325, 583

- Kauffmann, G., White, S. D. M., Heckman, T. M., Ménard, B., Brinchmann, J., Charlot, S., Tremonti, C., & Brinkmann, J. 2004, *MNRAS*, 353, 713
- Padmanabhan, N., et al. 2008, *ApJ*, 674, 1217
- Padmanabhan, N. et al. 2012, in preparation
- Pâris, I. et al. 2012, submitted to *Astronomy and Astrophysics*
- Pier, J. R., Munn, J. A., Hindsley, R. B., Hennessy, G. S., Kent, S. M., Lupton, R. H., & Ivezić, Ž. 2003, *AJ*, 125, 1559
- Predehl, P., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7732, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Ross, N. P., et al. 2012, *ApJS*, 199, 3
- Schlegel, D., White, M., & Eisenstein, D. 2009, in *ArXiv Astrophysics e-prints*, Vol. 2010, astro2010: The Astronomy and Astrophysics Decadal Survey, 314
- Schlegel, D. et al. 2012, in preparation
- Slosar, A., et al. 2011, *Journal of Cosmology and Astroparticle Physics*, 9, 1
- Smith, J. A., et al. 2002, *AJ*, 123, 2121
- Tegmark, M., et al. 2004, *ApJ*, 606, 702
- Tucker, D. L., et al. 2006, *Astronomische Nachrichten*, 327, 821
- Vanden Berk, D. E., et al. 2001, *AJ*, 122, 549
- Weinberg, D. H., Mortonson, M. J., Eisenstein, D. J., Hirata, C., Riess, A. G., & Rozo, E. 2012, *ArXiv e-prints*
- West, A. A., Hawley, S. L., Bochanski, J. J., Covey, K. R., Reid, I. N., Dhital, S., Hilton, E. J., & Masuda, M. 2008, *AJ*, 135, 785
- White, M., et al. 2011, *ApJ*, 728, 126
- Yanny, B., et al. 2009, *AJ*, 137, 4377
- York, D. G., et al. 2000, *AJ*, 120, 1579